arXiv:1612.01082v2 [cs.CV] 4 Jan 2017

Multi-Label Image Classification with Regional Latent Semantic Dependencies

Junjie Zhang*,1,3 Qi Wu*,2 Chunhua Shen[2] Jian Zhang[1] Jianfeng Lu[3]
[1]University of Technology Sydney, Australia [2]University of Adelaide, Australia
[3]Nanjing University of Science and Technology, China
{junjie.zhang@student.,jian.zhang@}uts.edu.au {qi.wu01,chunhua.shen}@adelaide.edu.au
lujf@njust.edu.cn

Abstract

Recent state-of-the-art approaches of multi-label image classification exploit the label dependencies at the global level, largely improving the labeling ability. However, accurately predicting small objects and visual concepts is still challenging due to the limited discrimination of the global visual features. In this paper, we propose a Regional Latent Semantic Dependencies model (RLSD) to address this problem. The utilized model includes a fully convolutional localization architecture to localize the regions that may contain multiple highly-dependent labels. The localized regions are further sent to the recurrent neural networks (RNN) to characterize the latent semantic dependencies at the regional level. Experimental results on several benchmark datasets show that our proposed model achieves the best performance compared to the state-of-the-art models, especially for predicting small objects occurred in the images. In addition, we set up an upper bound model (RLSD+ft-RPN) using bounding box coordinates during training, the experimental results also show that our RLSD can approach the upper bound without using the bounding-box annotations, which is more realistic in the real world.

1. Introduction

Image classification has been comprehensively studied for decades, especially for the single-label classification problem [19, 29, 41], as a traditional computer vision task, various progresses have been made on it. However, in the real world, an image usually contains abundant semantic information, such as objects, attributes, actions and scenes *etc*. With the available of the large-scale dataset and the enrichment of the data annotations, multi-label image classification has drawn lots of attentions [20, 21, 35]. Inspired by the advanced performance of convolutional neural networks (CNN) [23, 29, 43], various efforts have been made to apply the CNN on the multi-label classification problem [51, 52].

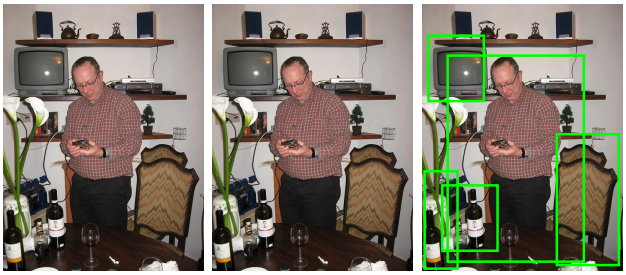

Ground truth: person; chair; dining table; bottle; wineglass; vase; TV

Multi-CNN: person; chair; TV

CNN+LSTM: person; chair; dining table; TV

RLSD: person; chair; dining table; bottle; wineglass; vase; TV

Figure 1: Example results of multi-label prediction from different models. The left is the ground-truth and the middle column shows the results from baseline models, Multi-CNN and CNN+LSTM. The right column displays the outputs of our proposed RLSD model, including predicted multiple labels and selected region proposals. Compare to the baseline methods, our model produces much richer predictions and is especially good at predicting small objects, such as 'bottle', 'wineglass' and 'vase' *etc*.

The most direct approach is to treat multi-label image classification problem as several separate single label classification problems, and to train the independent classifier for each label with a cross-entropy [20] or ranking loss (such as WARP [16]). Wei *et al*. [52] provide a regional solution allows predicting labels independently at the regional level. However, it is difficult for them to model the label dependencies between different labels. Intuitively, images with multiple labels usually contain strong correlations among the labels, for example, 'ocean' and 'ship' usually appear in the same image, while 'ocean' and 'cat' normally never occur together. To conveniently explore the label dependencies, the probabilistic graphical models (PGM) [49] are usually employed in the previous works [31, 46].

Most recently, Wang *et al*. [51] have shown that the recurrent neural networks (RNN) [24, 36] can efficiently capture the high order label dependencies. They unify the CNN and RNN as one framework to exploit the label dependencies at the global level, largely improving the labeling abil-

* Indicates equal contribution.

ity. However, predicting small objects and attributes is still challenging for these works due to the limited discrimination of the global visual features.

In this paper, our main contribution is that we propose a Regional Latent Semantic Dependencies (RLSD) model for the multi-label image classification, which effectively captures the latent semantic dependencies at the regional level. The proposed model combines the power of the region based features and the advantages of the RNN based label co-occurrence models, and achieves the best performance compared to the state-of-the-art multi-label classification models on several benchmark datasets, especially for predicting small objects and visual concepts. Fig. 1 shows an example output of our proposed RLSD model compared with the baselines models. We can see that the Multi-CNN and CNN+LSTM both fail to predict the 'bottle', 'vase' and 'wine glass' in the image due to their small size, while our model efficiently predicts them along with other large objects.

The framework of the proposed model is shown in Fig. 2. An input image is first processed through a CNN to extract convolutional features, which are further sent to an RPN-like (Regional Proposal Network) localization layer. Different from the conventional RPN in the object detection framework (such as faster R-CNN [40]) which tries to predict the proposals with a single object inside, our localization layer is designed to localize the regions in an image that may contain multiple semantically dependent labels. These regions are encoded with a fully-connected neural network and further sent to a RNN, which captures the latent semantic dependencies at the regional level. The RNN unit sequentially outputs a multi-class prediction, based on the outputs of the localization layer and the outputs of previous recurrent neurons. Finally, a max-pooling operation is carried out to fuse all the regional outputs as the final prediction. In addition, we also set up an upper bound model (RLSD+ft-RPN) using bounding-box coordinates for training. Our experimental results show that our model can approach the upper bound without using the bounding-box annotations, which is more realistic in the real world.

2. Related Works

Various works on multi-label image classification have been conducted over the past few years. The traditional bag-of-words (BoW) model [6, 7, 12, 22] relies on composing multiple modules, *e.g.*, feature representation (SIFT [34], HOG [11], LBP [37] *etc.*), classification (SVM [5], random forests [3]) and context modeling [6, 7, 22].

Recent progresses in the image classification are made based on the powerful deep convolutional neural networks, which try to model the high-level abstractions of the visual data by using architectures composed of multiple non-linear transformations. Several approaches have been proposed to expand single label classification network [23, 29, 43] to the multi-label problem. Gong *et al*. [16] combine top-k ranking objectives with a CNN architecture to tackle this problem. By defining a weight function for pair-wised ranking labels, they minimize the loss function so that positive labels are ranked higher than negative ones. Wei *et al*. [52] provide a regional solution that allows predicting labels independently at the regional level. They use BING [8] to generate object proposals and further send them into the CNN to compute multi-class scores. A max-pooling operation is applied to fuse the regional scores together as final classification results. We also use regional features and the max-pooling fusion. However, we consider the *regional latent semantic dependencies*, which allow us to *predict multiple labels jointly*.

There are also some works to solve the multi-label classification problem by designing a multi-modal representation, which bridges the semantic gap between the image and labels by learning the representation for the image visual content as well as labels. Canonical correlation analysis (CCA) [17] and kernel canonical correlation analysis (KCCA) [48] are usually carried out to build a latent semantic space to tackle the multi-label image annotation and retrieval problem. In addition, matrix completion [4] is also used to encode multi-label. These methods concentrate on digging labels' abundant semantic information, while neglect to explore the label dependencies among them.

To model the label dependencies, several approaches have been proposed. The probabilistic graphical models [49] are usually employed in previous works [2, 14, 31, 46, 54] to model the image feature-label joint distribution. There are several different graph structures to fulfill this purpose. For example, Chow-Liu tree [9] is used to build a tree based on the labels' mutual dependency in [2]. In [54], joint probability is exploited by directed acyclic graph and chain rules. Conditional random field is used in [14, 46] and matrix completion is used in [4]. A limitation of the graph-based approaches is that the richer the label semantic information is, the more complex the graph can be, which causes high computational complexity and low efficiency. Moreover, all of the above methods only model the label dependencies at the global level.

Recurrent Neural Network (RNN) has been proved to be able to model the temporal dependency effectively in a sequence and it has been successfully applied in several sequence-to-sequence problems, such as image captioning [26, 27, 50, 53], machine translation [45], speech recognition [18], language modeling [44], *etc*. Most recently, Wang *et al*. [51] have shown that the RNN [24, 36] can efficiently capture the high order label dependencies. They unify CNN and RNN into one framework to exploit the label dependencies at the global level, largely improving the labeling ability. We apply the RNN to capture the label dependencies as

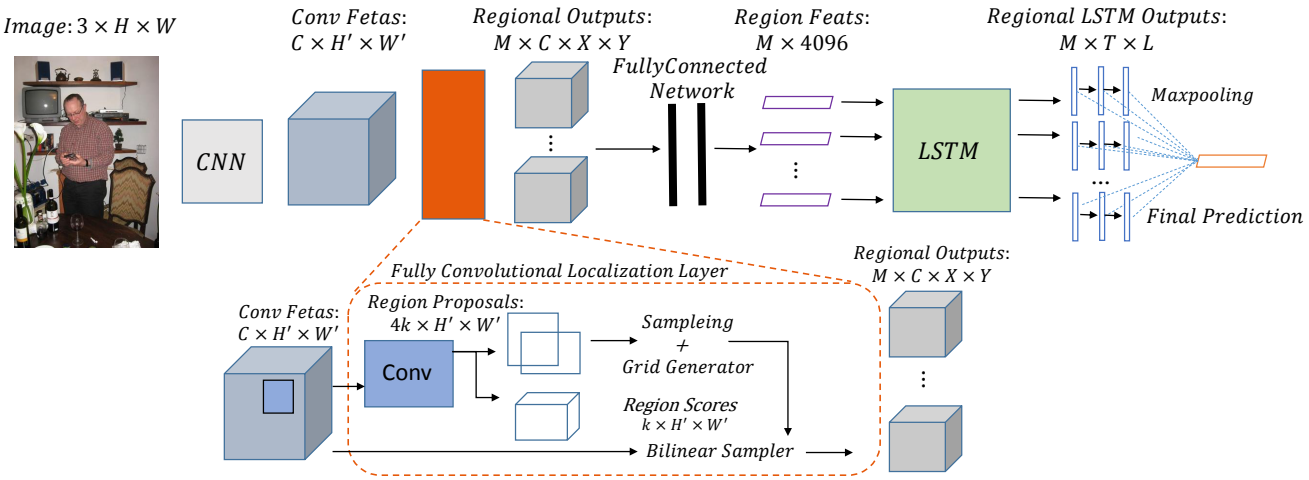
Image: 3 × H × W
Conv Fetas: C × H′ × W′
Regional Outputs: M × C × X × Y
FullyConnected Network
Region Feats: M × 4096
Regional LSTM Outputs: M × T × L
CNN
LSTM
Maxpooling
Final Prediction
Fully Convolutional Localization Layer
Regional Outputs: M × C × X × Y
Conv Fetas: C × H′ × W′
Region Proposals: 4k × H′ × W′
Conv
Sampleing + Grid Generator
Region Scores k × H′ × W′
Bilinear Sampler

Figure 2: Our proposed Regional Latent Semantic Dependencies model. An input image is first processed through a CNN to extract convolutional features, which are further sent to an RPN-like fully convolutional localization layer. The localization layer localize the regions in an image that potentially contain multiple highly-dependent labels. These regions are encoded with a fully-connected neural network and sent to the regional LSTM. Finally, a max-pooling operation is carried out to fuse all the regional outputs as the final prediction.

well. But different from [51], our regional latent semantic dependencies model considers the label dependencies at the *regional level*, which allows us to *predict small-size objects and visual concepts*.

In summary, our proposed model is inspired by the previous works in image classification [51, 52], object detection [40] and image captioning [26, 50]. We propose to utilize the region proposal network, fully-connected recognition network and RNN together to extract image regions with abundant semantic information, and explore the latent semantic dependencies simultaneously. The following sections describe details of our proposed RLSD model.

3. The RLSD Model

Overview. The key characteristic of our proposed model is that it can capture the regional semantic label dependencies. The novelty lies in the fact that this is achieved by a localization architecture, followed by a few of LSTMs (Long-Short Term Memory). The purpose of the localization layer is to localize the regions that contain multiple highly-dependent labels, while the LSTMs are employed to characterize the latent semantic label dependencies in a sequential manner. A max-pooling operation is carried out to finally fuse all the regional outputs. Fig. 2 shows the entire network of our proposed model.

In the following part, the localization layer is first introduced in Sec. 3.1 and the LSTM based label sequence prediction model is described in Sec. 3.2. The final max-pooling operation and the loss function is summarized in Sec. 3.3. The model initialization and some training details are given in Sec. 3.4.

3.1. Localizing Multi-label Regions

To explore the image at the regional level, we need to generate the regions that potentially contain the multiple objects and visual concepts. Hence, the first component of our proposed model is to localize these regions. The conventional object proposal algorithms (such as Selective Search [47], Objectness [1], BING [8] and MCG [39], *etc*.) are ruled out since these methods only focus on predicting single object proposals, which means that a proposed region normally only contains one single object. Instead, Johnson *et al*. [26] propose a fully convolutional neural network, extended from the Region Proposal Network (RPN) [40], to localize regions that can be described by a sentence, instead of a single label. Therefore, the proposed regions in [26] normally enjoy bigger label density, as well as the label complexity. Inspired by their work, we develop our approach of generating region proposals that are tailored for multi-label image classification.

3.1.1 Convolutional Features as Input

Since the convolutional layers of the CNN still preserve the spatial information of an image, which is necessary for us to explore the semantic dependencies at the regional level, we use them to extract image features. Specifically, we use the VGGNet [43] convolutional layer configuration, which consists of 13 convolutional layers with 3×3 kernel size and 5 max-pooling layers with 2×2 kernel size. The output of the last convolutional layer is used as image features. Given an input image with size $3 \times H \times W$, the convolutional features will be $C \times H' \times W'$, where $H' = \lfloor \frac{H}{16} \rfloor$, $W' = \lfloor \frac{W}{16} \rfloor$, and $C = 512$, same as in the VGGNet setting. The convolutional features are further sent to the localization layer to generate the region proposals that we are interested in.

3.1.2 Fully Convolutional Localization Layer

The input of the localization layer is the convolutional features extracted from the last step, while outputs are numbers of spatial regions of interest with a fixed-sized representation for each region.

Anchors and Regression. By referring to [26, 40], we predict region proposals by regressing offsets from a set of anchors. Specifically, each point inside the convolutional feature map is projected back into the original image ($H \times W$), and k different aspect ratio anchor boxes are generated at each projected point. Each anchor box is further sent into a fully convolutional network to produce the predicted box coordinates and a confidence score. The fully convolutional network consists of 256 convolutional filters with 3×3 kernel size, a ReLU layer and a final convolutional layer with $(4+1)k$ filters, where 4 stands for the number of the box coordinates, and 1 stands for the confidence score. We set $k = 12$ in our proposed models. As for the bounding-box regression from the anchors to the region proposals, we refer to [40] for the parameterization. More details can be found in the supplementary material.

Box Sampling and Bilinear Interpolation. A sampling mechanism is employed here to subsample the generated region proposals, since it is expensive to send all the proposals to the further LSTM-based label generation step. By referring to [26, 40], a $M = 256$ size minibatch is sampled. A box with IoU (Intersection of Union) which is larger than 0.7 ratio of the ground-truth region is considered as the positive samples, and less than 0.3 as the negative. We also restricted that at most of the half of boxes in one minibatch are positive samples and the other half are negative. During the test stage, non-maximum suppression is used to select the top M highest ranked proposals [1].

To ensure that the region proposal features can be accepted by the fully-connected layer and gradients can be back propagated to both the input features and box coordinates, the bilinear interpolation is used to replace the ROI pooling layer in [40]. We refer to the bilinear sampling operation in [26], which results in $M \times C \times X \times Y$ feature maps for the top M region proposals, where $C = 512$ is the VGGNet convolutional feature map size, and X, Y are the bilinear sampling grid size. In our case, we set $X = Y = 7$, referring to [25, 26].

3.1.3 Encoded by a Fully-Connected Network

After the regional features as $M \times C \times X \times Y$ are obtained, they are sent to a fully-connected network, which is formed by two 4096-d fully-connected layers and regularized by dropout. Features from each region are flattened into a vector and passed through this fully-connected network. Thus, each proposal region is encoded as a feature vector v with 4096 dimensions. All the regions' fully-connected features form a minibatch $V = [v_1, v_2, ...v_i, ...v_M]$ with the size of $M \times 4096$, where i indicates the i^{th} region proposal.

[1] In practical, the entire image is also added into the proposals, since some labels are related to the whole image.

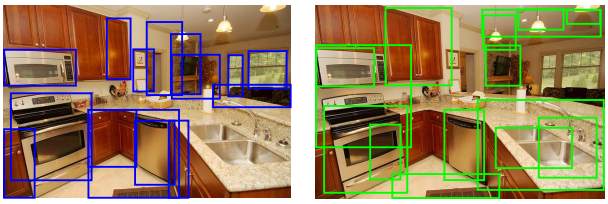

Figure 3: The comparison results between the Top-15 regions generated by MCG [39] (left) and our localization layer (right). Some of our generated regions contain multiple objects, for example, the bounding box on the most left side contains the object of 'oven', 'microwave' and 'kitchenwares' all together.

Fig. 3 shows an example of the comparison result between our localization layer proposed regions and the MCG [39] produced object proposals. The bounding boxes generated by our model are normally bigger, and some of them contain multiple objects. To show the effectiveness of the localization layer, we set a baseline model that uses the MCG [39] to replace our multi-label region localization layer for the further multi-label classification. More details can be found in the Sec. 4.

3.2. An LSTM-based Multi-Label Generator

To capture the latent semantic dependencies existed in those regions, we employ LSTMs to generate the sequence of label probability distributions on each single region.

The LSTM is a memory cell which encodes the knowledge at every time step for what inputs that have been observed up to this step. We follow the model used in [24]. Letting σ be the sigmoid nonlinearity, the LSTM updates for time step t given inputs x_t, h_{t-1}, c_{t-1} are:

$$
\begin{aligned}
i_t &= \sigma(W_{xi}x_t + W_{hi}h_{t-1} + b_i) && (1)\\
f_t &= \sigma(W_{xf}x_t + W_{hf}h_{t-1} + b_f) && (2)\\
o_t &= \sigma(W_{xo}x_t + W_{ho}h_{t-1} + b_o) && (3)\\
g_t &= \tanh(W_{xc}x_t + W_{hc}h_{t-1} + b_c) && (4)\\
c_t &= f_t \odot c_{t-1} + i_t \odot g_t && (5)\\
h_t &= o_t \odot \tanh(c_t) && (6)\\
p_t &= \mathrm{softmax}(h_t) && (7)
\end{aligned}
$$

Here, i_t, f_t, c_t, o_t are the input, forget, memory, output state of the LSTM. The various W matrices are trained parameters and $\odot$ represents the product with a gate value. h_t is the hidden state at time step t and is fed to a Softmax, which will produce a probability distribution p_t over all labels.

Given a region feature vector v, we set $x_0 = W_{ev}v$, where W_{ev} is the learnable region features embedding weights. Following the equations (1) to (6), it gives us an initial hidden state h_0 which can be used in the next time

step. From $t = 1$ to $t = T$, we set $x_t = W_{es}S_t$ and the hidden state h_{t-1} is given by the previous step, where W_{es} is the learnable label embedding weights. T is the numbers of label in a region and S_t is the input label at time step t [2]. Practically, in our RLSD model, since only the global multi-label ground-truth are provided and no regional ground-truth can be used in the training stage (as well as in the testing), we call the S_t as a latent label, which can be obtained by the following equation:

$$S_t^i = \mathbb{1}\{p_{t-1}^i = \max(p_{t-1})\} \tag{8}$$

where $\mathbb{1}$ is an indicator function and S_t is a one-hot vector that index $i = 1$, and 0 elsewhere. i is the index of the maximum value of the probability distribution p_{t-1} over all the labels, which is computed by the LSTM feed-forward process at the previous $t-1$ time step. After all the labels in a region are predicted, an 'END' label is added to finalize the prediction.

Giving all the M region features in a minibatch (all the regions in a minibatch come from the same image) into the LSTM model, we gather the prediction p_{tm} at each time step t, on each region m, to form a matrix with the shape $M \times T \times L$, where L is the label size of a dataset. If a region label length is less than T, we will pad 0.

3.3. Max-pooling and Loss Function

To suppress the possibly noisy prediction on some region proposals or at certain time steps, a cross region and time max-pooling is carried out to fuse the outputs into one integrative prediction. Suppose p_{tm} is the output prediction of region m at time step t and $p_{tm}^{(j)}(j = 1,...,L)$ is the j^{th} component of p_{tm}. The max-pooling in the fusion layer can be formulated as:

$$p^{(j)} = \max(p_{11}^{(j)},...,p_{1m}^{(j)},p_{21}^{(j)},...,p_{TM}^{(j)}) \tag{9}$$

where $p^{(j)}$ can be considered as the predicted value for the j^{th} category of the given image.

The max-pooling fusion is a crucial step for the proposed RLSD model to be robust to the noise. The output of the fusion layer is fed into a multi-way softmax layer with the squared loss as the cost function, which is defined as:

$$J = \frac{1}{N}\sum_{i=1}^{N}\sum_{l=1}^{L}(p_{il} - \hat{p}_{il})^2 \tag{10}$$

where $\hat{p}_i = y_i/||y_i||_1$ is the ground-truth probability vector of i^{th} image and p_i is the predictive probability vector of i^{th} image. N is the number of images.

[2] In practical, a label is represented as a one-hot vector, where 1 indicates the label exists, and 0 elsewhere.

3.4. Initialization and Pre-Training

Our model is capable of training end-to-end, but a proper initialization and pre-training mechanism is important to achieve a promising performance.

Localization Layer Pre-Training. The localization layer is pre-trained on the Visual Genome region caption dataset [28]. Different from other object detection datasets, each region in the image of this dataset normally contains several object and visual concepts, which is very suitable for our multi-label region localization task.

LSTM Pre-training. In the training stage, the LSTM is first pre-trained on the global image without region proposals, where every time step has the global image label as ground-truth to compute loss [3]. Then the pre-trained LSTM is used as the initialization of the regional LSTM in our proposed RLSD model. We find this initialization process is important for the model converging.

4. Experiments

In this section, we present our experimental results and analysis to show the effectiveness of our proposed RLSD model for multi-label image classification problem. We evaluate the proposed model on three benchmark datasets: VOC PASCAL 2007 [13], Microsoft COCO [32] and NUS-WIDE [10]. By comparing with several state-of-the-art models and baseline models, we show that our proposed RLSD model achieves the best performance. We further analyze the precision-recall along with the bounding-box size to show our model is especially good at predicting small objects.

4.1. Implementation Details

As mentioned in the Sec. 3.1, the VGGNet [43] is used to initialize the convolutional layers for the localization network. The embedding size of the image region feature is 512 (dimension of W_{ev}) and the label embedding size is 64 (dimension of W_{es}). We only use the one-layer LSTM, while the LSTM memory cell size is 512. Stochastic gradient descent (SGD) is used for optimization and we employ learning rate 10^{-5}, momentum 0.9 and drop rate 0.5.

Data preprocessing For VOC 2007 [13] and MS-COCO dataset [32], the ground-truth of bounding-boxes for the single objects are provided. Therefore, we can use the bounding-box coordinates to fine tune the localization layer in our model. By using this additional information, the fine-tuned model can extract the regions more accurately. Therefore, it can be viewed as an upper bound model to verify the RLSD's generality. We call this upper bound model

[3] This global model is also implemented and compared in the experimental part, known as CNN-LSTM

	plane	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	motor	person	plant	sheep	sofa	train	tv	mAP
INRIA [22]	77.2	69.3	56.2	66.6	45.5	68.1	83.4	53.6	58.3	51.1	62.2	45.2	78.4	69.7	86.1	52.4	54.4	54.3	75.8	62.1	63.5
FV [38]	75.7	64.8	52.8	70.6	30.0	64.1	77.5	55.5	55.6	41.8	56.3	41.7	76.3	64.4	82.7	28.3	39.7	56.6	79.7	51.5	58.3
CNN-SVM [42]	88.5	81.0	83.5	82.0	42.0	72.5	85.3	81.6	59.9	58.5	66.5	77.8	81.8	78.8	90.2	54.8	71.1	62.6	87.4	71.8	73.9
I-FT [52]	91.4	84.7	87.5	81.8	40.2	73.0	86.4	84.8	51.8	63.9	67.9	82.7	84.0	76.9	90.4	51.5	79.9	54.3	89.5	65.8	74.5
HCP-1000C [52]	95.1	90.1	92.8	89.9	51.5	80.0	91.7	91.6	57.7	77.8	70.9	89.3	89.3	85.2	93.0	64.0	85.7	62.7	94.4	78.3	81.5
HCP-2000C [52]	96.0	92.1	93.7	93.4	58.7	84.0	93.4	92.0	62.8	89.1	76.3	91.4	95.0	87.8	93.1	69.9	90.3	68.0	96.8	80.6	85.2
CNN-RNN [51]	96.7	83.1	**94.2**	92.8	61.2	82.1	89.1	94.2	64.2	83.6	70.0	92.4	91.7	84.2	93.7	59.8	93.2	75.3	**99.7**	78.6	84.0
Multi-CNN	95.8	88.7	93.2	91.1	51.5	82.1	89.5	91.8	68.3	80.2	75.7	91.4	92.6	88.7	92.8	61.2	82.9	68.0	96.2	78.0	83.0
CNN+LSTM	**96.8**	89.7	93.3	90.4	54.6	85.6	89.0	92.3	68.9	80.9	**76.8**	90.5	93.6	89.5	93.0	60.3	84.3	67.4	96.2	76.7	83.5
MCG-CNN+LSTM	94.3	90.4	91.4	91.4	64.9	90.3	92.9	94.3	71.9	83.7	55.6	92.5	94.5	90.0	96.8	60.1	90.5	71.3	96.3	84.7	84.9
RLSD	95.3	92.4	91.2	92.1	**71.9**	91.1	93.3	94.8	**74.9**	86.1	70.4	93.3	95.6	89.7	**98.0**	66.8	89.4	**75.7**	96.6	85.9	87.3
RLSD+ft-RPN	96.4	**92.7**	93.8	**94.1**	71.2	**92.5**	**94.2**	**95.7**	74.3	**90.0**	74.2	**95.4**	**96.2**	**92.1**	97.9	**66.9**	**93.5**	73.7	97.5	**87.6**	**88.5**

Table 1: Comparisons of classification results (AP in %) on the PASCAL VOC 2007 dataset.

RLSD+ft-RPN (RLSD with fine-tuned RPN). The data pre-processing procedures are: given an image with several single-object bounding-boxes, we first compute their centers. Since the relevant objects usually appear close to each other, Euclidean distance based hierarchical clustering is applied on these bounding-box centers to generate several clusters. Bounding-boxes which belong to the same cluster are merged as one novel region. These novel regions are further used as the ground-truth to fine-tune the localization layer. Please note that only the bounding-box coordinates are used in this pre-processing, there are no label annotations involved.

4.2. Evaluation Metrics

We refer to the evaluation metrics used in [16, 51] to compute precision and recall for the predicted labels. For each test image, we predict k highest ranked labels and compare against the image ground-truth. The precision is the number of the correctly annotated labels divided by the number of predicted labels; the recall is the number of correctly annotated labels divided by the number of ground-truth labels. We compute *overall precision & recall (op & or)* and *per-class precision & recall (cp & cr)* based on the formulas blew. We also compute the *mean average precision* (mAP) for the comparisons.

$$op = \frac{\sum_{i=1}^{c} N_i^c}{\sum_{i=1}^{c} N_i^p}, or = \frac{\sum_{i=1}^{c} N_i^c}{\sum_{i=1}^{c} N_i^g}, \tag{11}$$

$$cp = \frac{1}{c}\sum_{i=1}^{c} \frac{N_i^c}{N_i^p}, cr = \frac{1}{c}\sum_{i=1}^{c} \frac{N_i^c}{N_i^g}, \tag{12}$$

where c is the number of total labels, N_i^c is the number of images that are accurately labeled for i_{th} label,N_i^p is the number of the predicted images for i_{th} label, and N_i^g is the number of the ground truth images for i_{th} label.

4.3. Baseline Models

To evaluate the effectiveness of our proposed RLSD model, we implement three baseline models for comparisons. The experimental results of these baseline models indicate the significance of our model's components.

Multi-label CNN: This is a standard CNN model without considering any label-dependencies. We use the VGGNet [43] pre-trained on the ImageNet for parameters initialization, and fine-tune the model on the benchmark datasets. The dimension of the last fully-connected layer is changed to the class number of each dataset and the element-wise logistic loss is applied. The learning rates of the last two fully-connected layers are initialized as 0.001 and 0.05 respectively, the rest of the layers are fixed. We run around 45 epochs in total until the model leads to convergence and decease the learning rate to one tenth every 15 epochs, the weight decay is 0.0005 and momentum is 0.9.

CNN+LSTM: This is a model with the same configurations as our proposed RLSD model, except that it only considers the label dependencies at the global level. We use the last fully-connected layer of the pre-trained VGGNet [43] as the global image representation. Then we feed the image feature into the LSTM to train a multi-label classification model. This baseline is similar to the CNN-RNN model proposed in [43], except the image feature is only fed into the LSTM once at the first time step. The dimensions of label embedding and LSTM layer are the same as our proposed model: 64 and 512 respectively.

MCG-CNN+LSTM: This is another baseline model, which is used to verify the effectiveness of our localization of multi-label regions. In this model, we use the exactly same configurations as our proposed model in Sec. 3 except that the region proposal network is replaced by an object proposal tool: the Multiscale Combinatorial Grouping (MCG) [39]. The pre-extracted top-256 object proposals (based on the object proposal confidence score) for each image are sent into our region-based LSTM for training and testing. As shown in Fig. 3, our region-based localization layer can generate the proposals with multiple labels.

4.4. Results on the VOC PASCAL 2007

VOC PASCAL 2007 [13] is a benchmark dataset for image classification, segmentation and detection. In VOC 2007 dataset, there are 9,963 images in total, of which 5,011 and 4,952 images for training/validation and testing respectively. 20 common objects are annotated in this dataset.

	C-P	C-R	C-F1	O-P	O-R	O-F1
Multi-CNN	47.5	85.2	61.0	46.4	90.2	61.3
CNN+LSTM	46.7	88.9	61.2	47.0	91.5	62.1
MCG-CNN+LSTM	48.6	88.7	62.8	46.9	91.4	62.0
RLSD	50.5	90.6	64.9	47.5	92.4	62.7
RLSD+ft-RPN	**51.6**	**90.7**	**65.8**	**47.6**	**92.7**	**62.9**

Table 2: Comparisons on the PASCAL VOC 2007 for different evaluation metrics, $k = 3$.

Tab. 1 shows the comparisons to the state-of-the-art methods including the HCP-1000C/2000C [52] and CNN-RNN [51], along with the baseline methods on VOC 2007 dataset. The HCP-1000C uses image region as the input and fine-tunes the CNN pre-trained on ImageNet 1000 classes. The HCP-2000C augments the ILSVRC-2012 training set with additional 1,000 ImageNet classes to pre-train the model. The CNN-RNN model [51], as we mentioned before, explores the image multi-label dependencies at the global level. Our proposed RLSD model uses the pre-trained region proposal network and fixes the localization layer during the training, while the upper bound mode RLSD+ft-RPN fine-tunes the localization layer with the region-based bounding-boxes as the guidance. The region bounding-boxes generation details have been discussed in the Sec. 4.1. From Tab. 1, we can see that our RLSD model outperforms those state-of-the-art methods and baseline models in most classes. It achieves the 87.3% mAP, while the CNN-RNN model [51] is 84% and our own CNN+LSTM baseline is 83.5%. This proves the regional semantic dependencies are more advanced than only considering the label dependencies at the global level. The gap between the Multi-CNN and our RLSD (4.3 %) is larger than the gap between Multi-CNN and CNN+LSTM (0.5%), which also stresses this point. Although the MCG-CNN+LSTM model uses the regional information, the region proposal from the MCG [39] normally only contains a single object, which means the label dependencies are invalid on those proposals anymore. In our RLSD model, the region-based localization layer can localize the regions that contain multiple highly-dependent labels. Therefore, our RLSD model also outperforms the MCG-CNN+LSTM model with a big margin, 87.3% VS. 84.9%, which is only 1.2% lower than the upper bound model RLSD+ft-RPN.

By introducing the region information into our model, it surpasses the current methods in a big gap for predicting small objects and visual concepts such as 'bottle', 'TV' and 'chair' *etc*. Take the 'bottle' as an example, we achieve 10% higher than CNN-RNN model. Tab. 2 compares the per-class precision & recall (cp & cr), overall precision & recall (op & or), and F1 score between our proposed models and the baseline models. Our RLSD outperforms the baselines on all the evaluation metrics, and approaches the RLSD+ft-RPN results. Particularly, our model has achieved

	C-P	C-R	C-F1	O-P	O-R	O-F1	mAP	mAP@10
WARP [16]	59.3	52.5	55.7	59.8	61.4	60.7	-	49.2
CNN-RNN [51]	66.0	55.6	60.4	69.2	**66.4**	**67.8**	-	61.2
Multi-CNN	54.8	51.4	53.1	56.7	58.6	57.6	60.4	57.8
CNN+LSTM	62.1	51.2	56.1	68.1	56.6	61.8	61.8	59.2
MCG-CNN+LSTM	64.2	53.1	58.1	61.3	59.3	61.3	64.4	62.4
RLSD	**67.7**	56.4	61.5	**70.5**	59.9	64.8	67.4	65.9
RLSD+ft-RPN	67.6	**57.2**	**62.0**	70.1	63.4	66.5	**68.2**	**66.7**

Table 3: Comparisons on the MS-COCO dataset for $k = 3$. mAP@10 measures are additionally computed for comparison.

a much higher precision, no matter for per-class or overall evaluations, and is very close to the upper bound results.

4.5. Results on the Microsoft COCO

Microsoft COCO dataset [32] is a large scale benchmark dataset for several vision tasks. There are in total 123,287 images for training and validation with 80 object concepts annotated. We use all the annotated object labels in an image as the multi-label ground-truth, and employ its training set as the training data and validation set as the test data. After removing the images without annotation, we have 82,081 images for training and 40,137 images for testing. We obtain the semantic dependencies of these labels by computing their co-occurrence rates and form them as a matrix (see our supplementary material for the visualization of co-occurrence matrix). We have found that there are strong dependencies in its label set, such as 'keyboard' and 'computer' always appear together.

We compare our RLSD model with the WARP [16] model and CNN-RNN [51] model in Tab. 3. Our RLSD model outperforms the CNN-RNN [51] on the evaluation metrics of per-class precision & recall, and overall precision. Although we have a lower overall recall, because our model may predict less than k labels for an image (a threshold $t = 0.5$ is set on the prediction score to decide whether the label should be output). We have higher mAP and mAP@10 measures than the state-of-the-art methods. We also outperform the MCG-CNN+LSTM model, which only uses the MCG object proposals as the input. The RLSD+ft-RPN results shows , our model is still close to the upper bound even without the RPN fine-tuning, which indicates

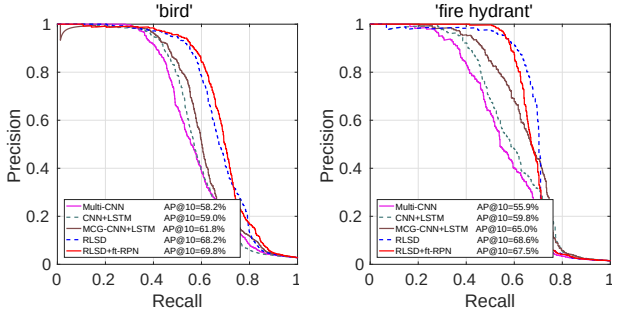

Figure 4: Precision-Recall curves for the 'bird' and 'fire hydrant' classes in the MS COCO dataset, for our RLSD models and the baseline models. The average precision @ 10 are also given in the figure.

person,baseball,baseball bat | person,car,traffic,horse,tie | cup,fork,knife,pizza | bottle,tv,mouse,keyboard,book

Figure 5: Some example multi-label classification results from the MS COCO, see more in the supplementary material

the generality of our model in the real world. Fig. 4 shows the precision-recall curves of two selected classes (bird and fire hydrant) in the MS COCO dataset, these two kinds of objects are visually small in the images of this dataset. The PR curves illustrate that our RLSD model significantly surpasses the baseline models and is especially good at predicting these small objects (All the PR curves of 80 classes in the MS COCO can be found in the supplementary material). Some qualitative results are shown in Fig. 5.

Fig. 6 analyzes the recall along with the object ground-truth bounding-box size on the MS COCO dataset, which shows that our model achieves much higher recall than the CNN+LSTM model on those labels that corresponding bounding-boxes are smaller. This means that our RLSD model is more sensitive to 'small' objects. Both recall curves start to fall when the object is so large that it almost fill the whole image, similar case is also observed in [51].

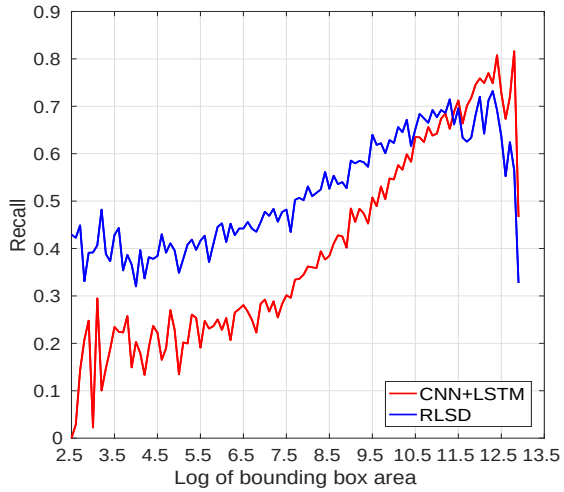

Figure 6: The relationship between recall and bounding-box area on Microsoft COCO dataset.

4.6. Results on the NUS-WIDE

NUS-WIDE [10] is a web image dataset associated with user tags. The whole dataset contains 269,648 images and 1,000 tags. These images are further manually labeled into 81 concepts. After removing the no-annotated images, the training set and test set contains 125,449 and 83,898 images respectively.

	C-P	C-R	C-F1	O-P	O-R	O-F1	mAP
Metric Learning [30]	-	-	-	-	-	21.3	-
Multi-edge graph [33]	-	-	-	35.0	37.0	36.0	-
KNN [10]	32.6	19.3	24.3	42.9	53.4	47.6	-
softmax [51]	31.7	31.2	31.4	47.8	59.5	53.0	-
WARP [16]	31.7	35.6	33.5	48.6	60.5	53.9	-
CNN-RNN [51]	40.5	30.4	34.7	49.9	61.7	55.2	-
Multi-CNN	40.2	42.8	41.5	53.4	66.4	59.2	47.1
CNN+LSTM	42.2	45.8	43.9	53.4	66.5	59.3	49.9
MCG-CNN+LSTM	44.3	48.1	46.1	54.1	67.2	59.9	52.4
RLSD	**44.4**	**49.6**	**46.9**	**54.4**	**67.6**	**60.3**	**54.1**

Table 4: Comparisons on the NUS-WIDE on 81 concepts for $k = 3$.

Tab. 4 shows the comparison results on the NUS-WIDE dataset on 81 concepts. We compare the proposed RLSD model with previous mentioned methods including metric learning [30], multi-edge graph [33], K nearest neighbor [10], softmax prediction, WARP method [16] and the latest state-of-the-art method CNN-RNN [51]. Same as [51], we do not fine-tune the CNN image representation. Our proposed RLSD model outperforms these methods in a large margin. Specifically, we achieve 4% higher precision than the CNN-RNN. Because the bounding-boxes are not provided in the NUS-WIDE dataset, there is no upper bound for this dataset.

5. Conclusion and Future Work

Multi-label image classification is an important problem on computer vision, because it is not only more challenging than the single-label image classification, but also closer to the real world applications. In this paper, we propose a Regional Latent Semantic Dependencies (RLSD) model to address this problem. As its name suggests, this proposed model can capture the label dependencies at the regional level. Experimental results on several benchmark datasets show that the proposed RLSD model consistently achieves the superior overall performance to the state-of-the-art approaches, especially for predicting small objects and visual concepts occurred in the image.

A limitation of our current model is that it may fail to localize the multi-label regions in some rarely encountered scenes. Although this can be addressed by fine-turning

the localization layer with pre-processed object bounding-boxes (such as our RLSD+ft-RPN model), it is not realistic when the human annotated object bounding-boxes are not available (such as the NUS-WIDE). The attention mechanism might be valued in this situation since it allows the label prediction process to attend on regions in an unsupervised manner. We will investigate that in our future work.

References

[1] B. Alexe, T. Deselaers, and V. Ferrari. Measuring the objectness of image windows. *IEEE Trans. Pattern Anal. Mach. Intell.*, 34(11):2189–2202, 2012. 3

[2] J. K. Bradley and C. Guestrin. Learning tree conditional random fields. In *Proc. Int. Conf. Mach. Learn.*, pages 127–134, 2010. 2

[3] L. Breiman. Random forests. *Mach. Learn.*, 45(1):5–32, 2001. 2

[4] R. S. Cabral, F. De la Torre, J. P. Costeira, and A. Bernardino. Matrix completion for multi-label image classification. In *Proc. Advances in Neural Inf. Process. Syst.*, volume 201, page 2, 2011. 2

[5] C.-C. Chang and C.-J. Lin. Libsvm: a library for support vector machines. *ACM Trans. Intelligent Syst. and Technology*, 2(3):27, 2011. 2

[6] Q. Chen, Z. Song, J. Dong, Z. Huang, Y. Hua, and S. Yan. Contextualizing object detection and classification. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(1):13–27, 2015. 2

[7] Q. Chen, Z. Song, Y. Hua, Z. Huang, and S. Yan. Hierarchical matching with side information for image classification. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 3426–3433. IEEE, 2012. 2

[8] M.-M. Cheng, Z. Zhang, W.-Y. Lin, and P. Torr. Bing: Binarized normed gradients for objectness estimation at 300fps. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 3286–3293, 2014. 2, 3

[9] C. Chow and C. Liu. Approximating discrete probability distributions with dependence trees. *IEEE Trans. Inf. Theory*, 14(3):462–467, 1968. 2

[10] T.-S. Chua, J. Tang, R. Hong, H. Li, Z. Luo, and Y.-T. Zheng. Nus-wide: A real-world web image database from national university of singapore. In *Proc. ACM Conf. Image and Video Retrieval*, Santorini, Greece., 2009. 5, 8

[11] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, volume 1, pages 886–893. IEEE, 2005. 2

[12] J. Dong, W. Xia, Q. Chen, J. Feng, Z. Huang, and S. Yan. Subcategory-aware object classification. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 827–834, 2013. 2

[13] M. Everingham, S. M. A. Eslami, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes challenge: A retrospective. *Int. J. Comput. Vision*, 111(1):98–136, Jan. 2015. 5, 6

[14] N. Ghamrawi and A. McCallum. Collective multi-label classification. In *Proc. ACM Conf. Information and Knowledge Management*, pages 195–200. ACM, 2005. 2

[15] R. Girshick. Fast r-cnn. In *Proc. IEEE Int. Conf. Comp. Vis.*, pages 1440–1448, 2015. 10

[16] Y. Gong, Y. Jia, T. Leung, A. Toshev, and S. Ioffe. Deep convolutional ranking for multilabel image annotation. *arXiv preprint arXiv:1312.4894*, 2013. 1, 2, 6, 7, 8

[17] Y. Gong, Q. Ke, M. Isard, and S. Lazebnik. A multi-view embedding space for modeling internet images, tags, and their semantics. *Int. J. Comput. Vision*, 106(2):210–233, 2014. 2

[18] A. Graves, A.-r. Mohamed, and G. Hinton. Speech recognition with deep recurrent neural networks. In *Proc. IEEE Conf. Acoustics, Speech and Signal Processing*, pages 6645–6649. IEEE, 2013. 2

[19] G. Griffin, A. Holub, and P. Perona. Caltech-256 object category dataset. 2007. 1

[20] M. Guillaumin, T. Mensink, J. Verbeek, and C. Schmid. Tagprop: Discriminative metric learning in nearest neighbor models for image auto-annotation. In *Proc. IEEE Int. Conf. Comp. Vis.*, pages 309–316. IEEE, 2009. 1

[21] Y. Guo and S. Gu. Multi-label classification using conditional dependency networks. In *Proc. Int. Joint Conf. Artificial Intell.*, volume 22, page 1300, 2011. 1

[22] H. Harzallah, F. Jurie, and C. Schmid. Combining efficient object localization and image classification. In *Proc. IEEE Int. Conf. Comp. Vis.*, pages 237–244. IEEE, 2009. 2, 6

[23] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015. 1, 2

[24] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997. 1, 2, 4

[25] M. Jaderberg, K. Simonyan, A. Zisserman, et al. Spatial transformer networks. In *Proc. Advances in Neural Inf. Process. Syst.*, pages 2017–2025, 2015. 4

[26] J. Johnson, A. Karpathy, and L. Fei-Fei. Densecap: Fully convolutional localization networks for dense captioning. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2016. 2, 3, 4, 10

[27] A. Karpathy and L. Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 3128–3137, 2015. 2

[28] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D. A. Shamma, M. Bernstein, and L. Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. 2016. 5

[29] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Proc. Advances in Neural Inf. Process. Syst.*, pages 1097–1105, 2012. 1, 2

[30] J. Li, X. Lin, X. Rui, Y. Rui, and D. Tao. A distributed approach toward discriminative distance metric learning. *IEEE Trans. Neural Netw. & Learn. Syst.*, 26(9):2111–2122, 2015. 8

[31] X. Li, F. Zhao, and Y. Guo. Multi-label image classification with a probabilistic label enhancement model. *Proc. Uncertainty in Artificial Intell*, 2014. 1, 2

[32] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *Proc. Eur. Conf. Comp. Vis.*, pages 740–755. Springer, 2014. 5, 7

[33] D. Liu, S. Yan, Y. Rui, and H.-J. Zhang. Unified tag analysis with multi-edge graph. In *Proc. ACM Conf. on Multimedia*, pages 25–34. ACM, 2010. 8

[34] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision*, 60(2):91–110, 2004. 2

[35] A. Makadia, V. Pavlovic, and S. Kumar. A new baseline for image annotation. In *Proc. Eur. Conf. Comp. Vis.*, pages 316–329. Springer, 2008. 1

[36] T. Mikolov, M. Karafiát, L. Burget, J. Cernockỳ, and S. Khudanpur. Recurrent neural network based language model. In *Interspeech*, volume 2, page 3, 2010. 1, 2

[37] T. Ojala, M. Pietikäinen, and D. Harwood. A comparative study of texture measures with classification based on featured distributions. *Pattern Recogn.*, 29(1):51–59, 1996. 2

[38] F. Perronnin, J. Sánchez, and T. Mensink. Improving the fisher kernel for large-scale image classification. In *Proc. Eur. Conf. Comp. Vis.*, pages 143–156. Springer, 2010. 6

[39] J. Pont-Tuset, P. Arbeláez, J. Barron, F. Marques, and J. Malik. Multiscale combinatorial grouping for image segmentation and object proposal generation. In *arXiv preprint arXiv:1503.00848*, March 2015. 3, 4, 6, 7

[40] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Proc. Advances in Neural Inf. Process. Syst.*, pages 91–99, 2015. 2, 3, 4

[41] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. Imagenet large scale visual recognition challenge. *Int. J. Comput. Vision*, 115(3):211–252, 2015. 1

[42] A. Sharif Razavian, H. Azizpour, J. Sullivan, and S. Carlsson. Cnn features off-the-shelf: an astounding baseline for recognition. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn. Workshop*, pages 806–813, 2014. 6

[43] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. 2015. 1, 2, 3, 5, 6

[44] M. Sundermeyer, R. Schlüter, and H. Ney. Lstm neural networks for language modeling. In *Interspeech*, pages 194–197, 2012. 2

[45] I. Sutskever, O. Vinyals, and Q. V. Le. Sequence to sequence learning with neural networks. In *Proc. Advances in Neural Inf. Process. Syst.*, 2014. 2

[46] M. Tan, Q. Shi, A. van den Hengel, C. Shen, J. Gao, F. Hu, and Z. Zhang. Learning graph structure for multi-label image classification via clique generation. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 4100–4109, 2015. 1, 2

[47] J. R. R. Uijlings, K. E. A. van de Sande, T. Gevers, and A. W. M. Smeulders. Selective search for object recognition. *Int. J. Comput. Vision*, 104(2):154–171, 2013. 3

[48] T. Uricchio, L. Ballan, L. Seidenari, and A. Del Bimbo. Automatic image annotation via label transfer in the semantic space. *arXiv preprint arXiv:1605.04770*, 2016. 2

[49] L. C. van der Gaag and A. J. Feelders. Probabilistic graphical models. 1, 2

[50] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: A neural image caption generator. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 3156–3164, 2015. 2, 3

[51] J. Wang, Y. Yang, J. Mao, Z. Huang, C. Huang, and W. Xu. Cnn-rnn: A unified framework for multi-label image classification. *arXiv preprint arXiv:1604.04573*, 2016. 1, 2, 3, 6, 7, 8

[52] Y. Wei, W. Xia, J. Huang, B. Ni, J. Dong, Y. Zhao, and S. Yan. Cnn: Single-label to multi-label. *arXiv preprint arXiv:1406.5726*, 2014. 1, 2, 3, 6, 7

[53] Q. Wu, C. Shen, L. Liu, A. Dick, and A. van den Hengel. What value do explicit high level concepts have in vision to language problems? In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, June 2016. 2

[54] M.-L. Zhang and K. Zhang. Multi-label learning by exploiting label dependency. In *Proc. ACM SIGKDD Int. Conf. Knowledge discovery & data mining*, pages 999–1008. ACM, 2010. 2

6. Supplementary Materials

This supplementary material includes following parts: at first, the details of the box regression in the localization layer is introduced in Sec. 6.1, and the visualization of co-occurrence matrix is described in Sec. 6.2. All the PR curves of 80 classes in the MS COCO is shown in Sec. 6.3. More example multi-label classification results from the MS COCO can be found in Sec. 6.4.

6.1. Box Regression in Localization Layer

By referring to [15, 26], we apply the log-space scaling transformations on the anchor box, which means given an anchor box's parameters (a_x,a_y,a_w,a_h), where (a_x,a_y) is the center of the anchor box, and a_w, a_h stands for the width and height of the anchor box respectively, we generate the region coordinates $b = (b_x,b_y,b_w,b_h)$ by following folumas:

$$b_x = a_x + t_x a_w \qquad b_y = a_y + t_y a_h \tag{13}$$

$$b_w = a_w exp(t_w) \qquad b_h = a_h exp(t_h) \tag{14}$$

The scalars t_x,t_y,t_w,t_h are predicted by our model. Smooth L_1 norm is employed as the loss function to regress the region location. Given the ground-truth coordinates $g = (g_x,g_y,g_w,g_h)$, the loss function is defined as:

$$L(b,g) = \sum_{i \in x,y,w,h} Smooth_{L_1}(b_i,v_i) \tag{15}$$

where

$$Smooth_{L_1} = \begin{cases} 0.5x^2 & if|x| < 1 \\ |x| - 0.5 & otherwise \end{cases} \tag{16}$$

6.2. Co-occurrence Matrix of MS COCO

We obtain the dependencies of MS COCO labels by computing their co-occurrence rates and form them as a matrix. Fig. 7 is the visualization of co-occurrence matrix.

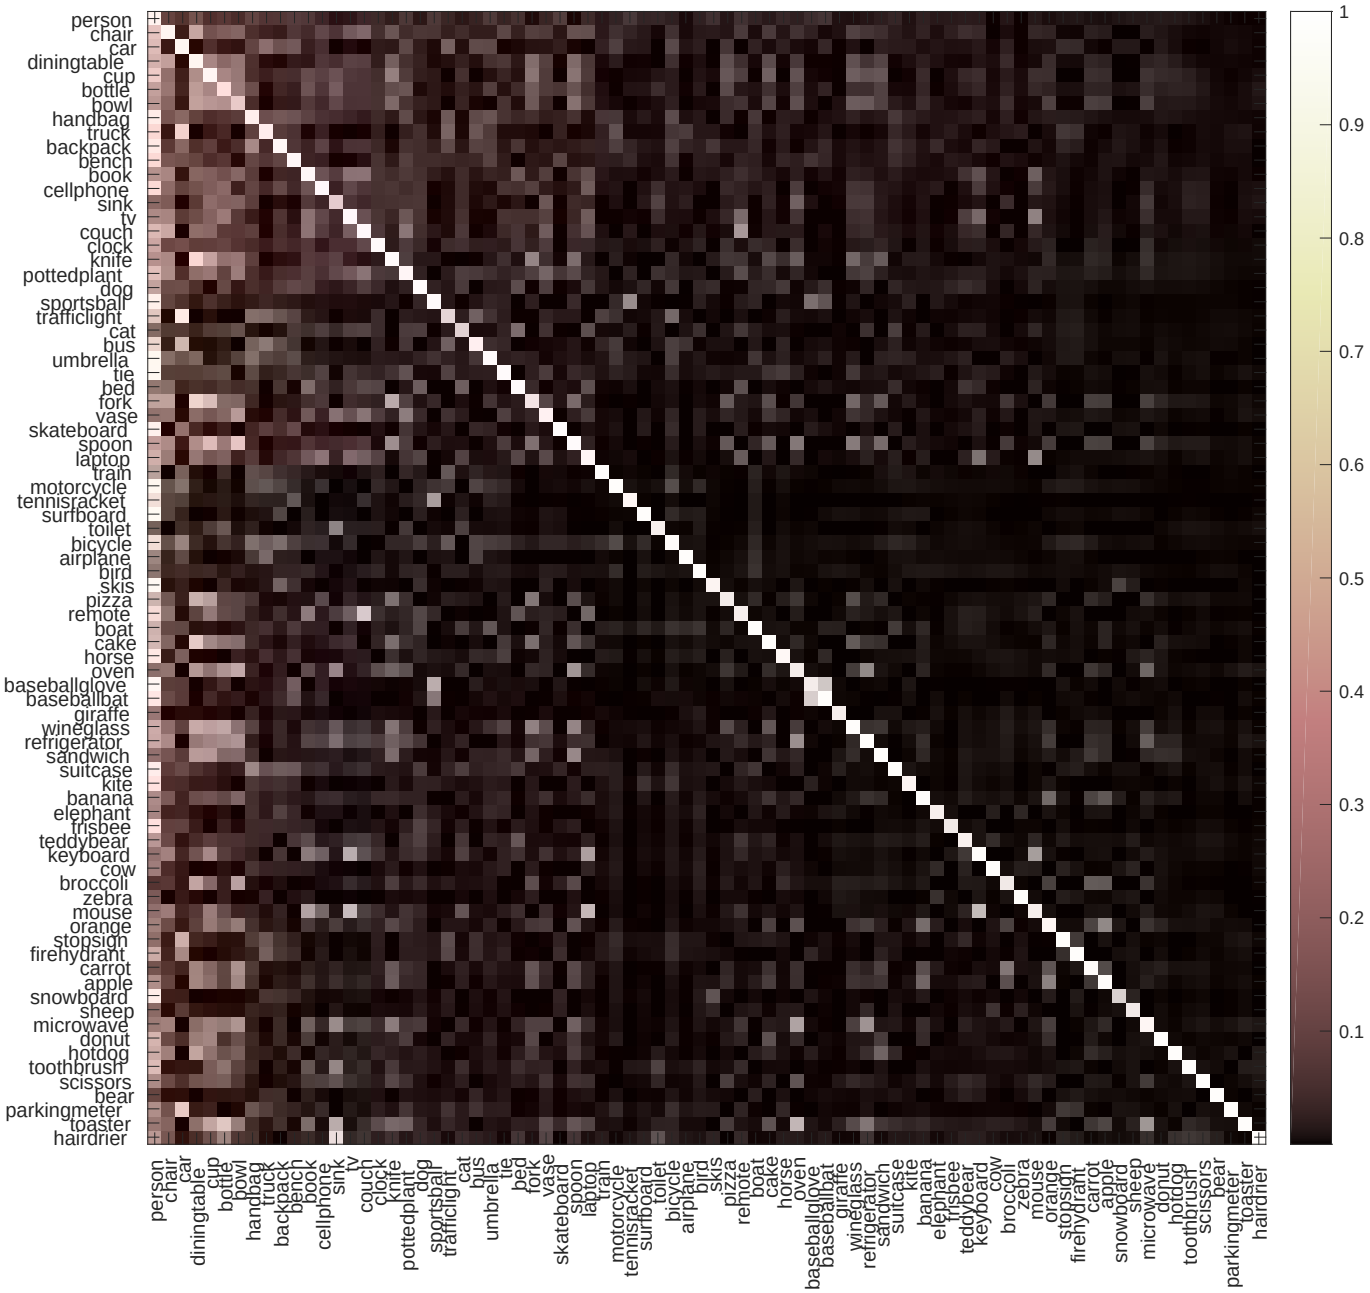

Figure 7: The visualization of COCO label co-occurrence matrix

6.3. PR Curves of MS COCO

All the PR curves of 80 classes in the MS COCO are shown in Fig.2 - 10.

6.4. Example Results of RLSD

Fig. 17 shows some qualitative results of our RLSD's prediction.

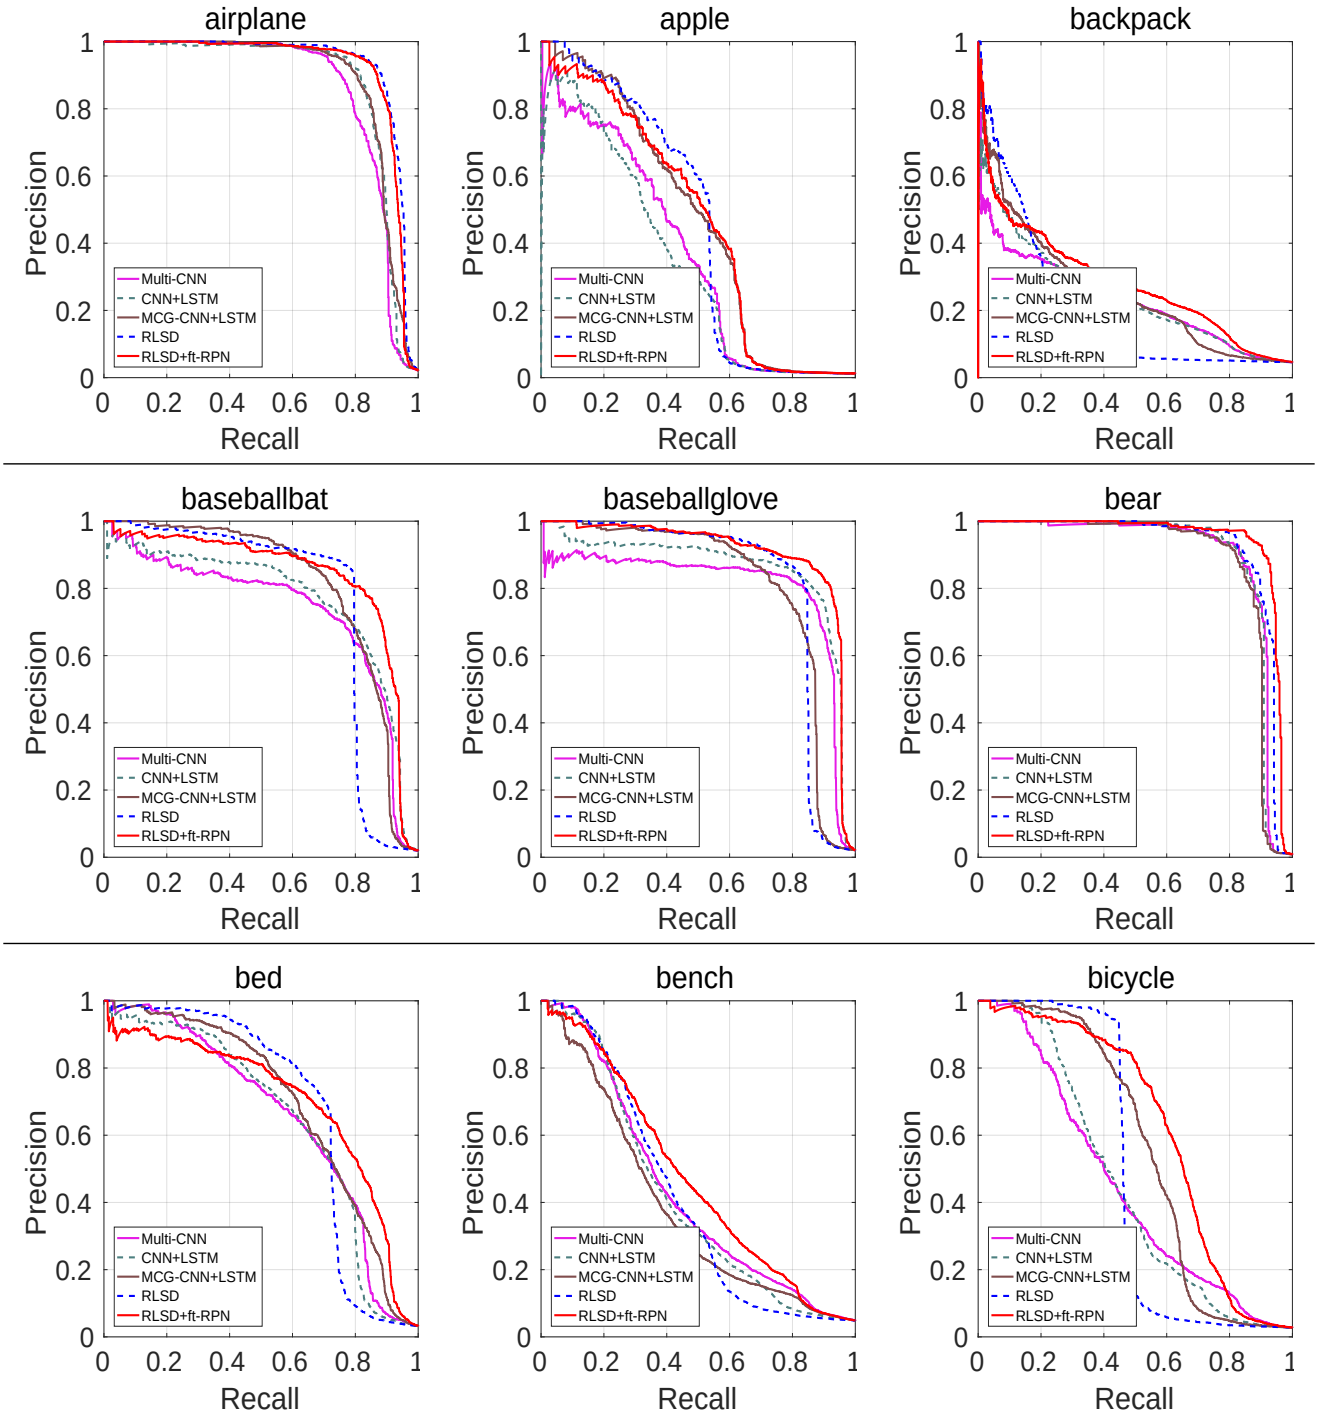

Figure 8: Precision-Recall curves for all 1st-9th classes in the MS COCO dataset, for our RLSD models and the baseline models.

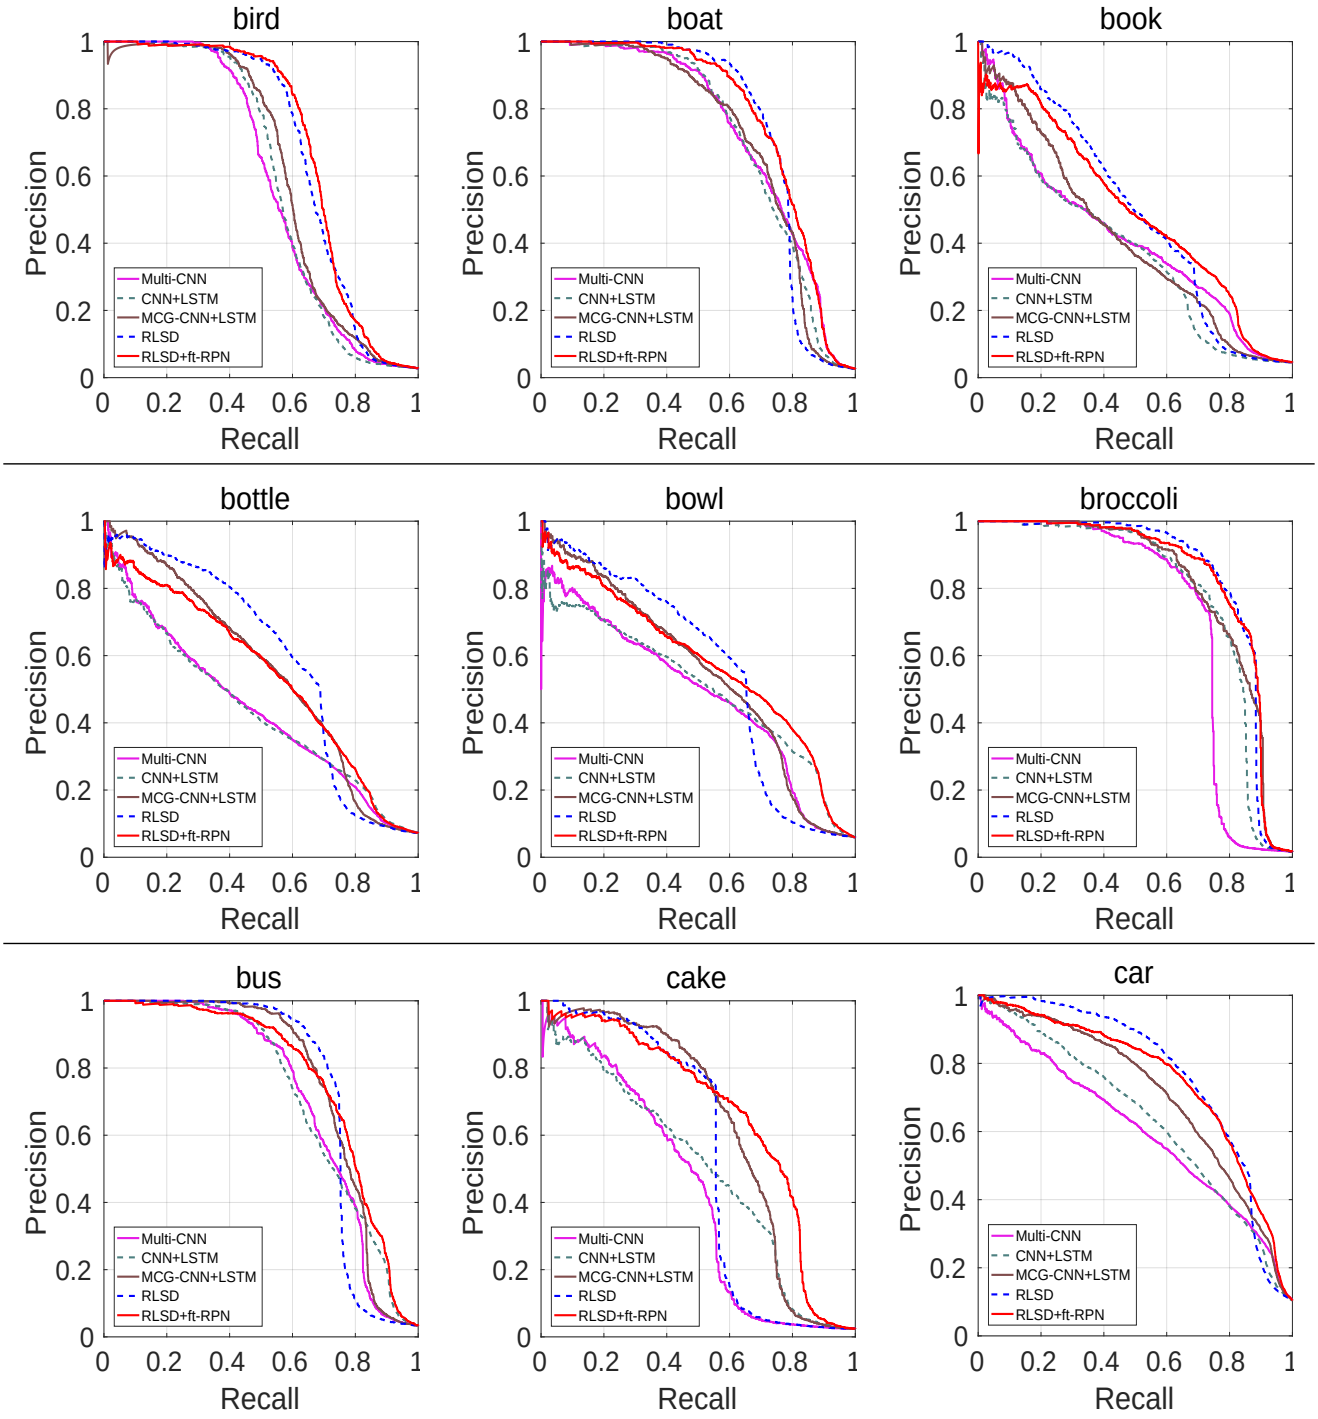

Figure 9: Precision-Recall curves for all 10th-18th classes in the MS COCO dataset, for our RLSD models and the baseline models.

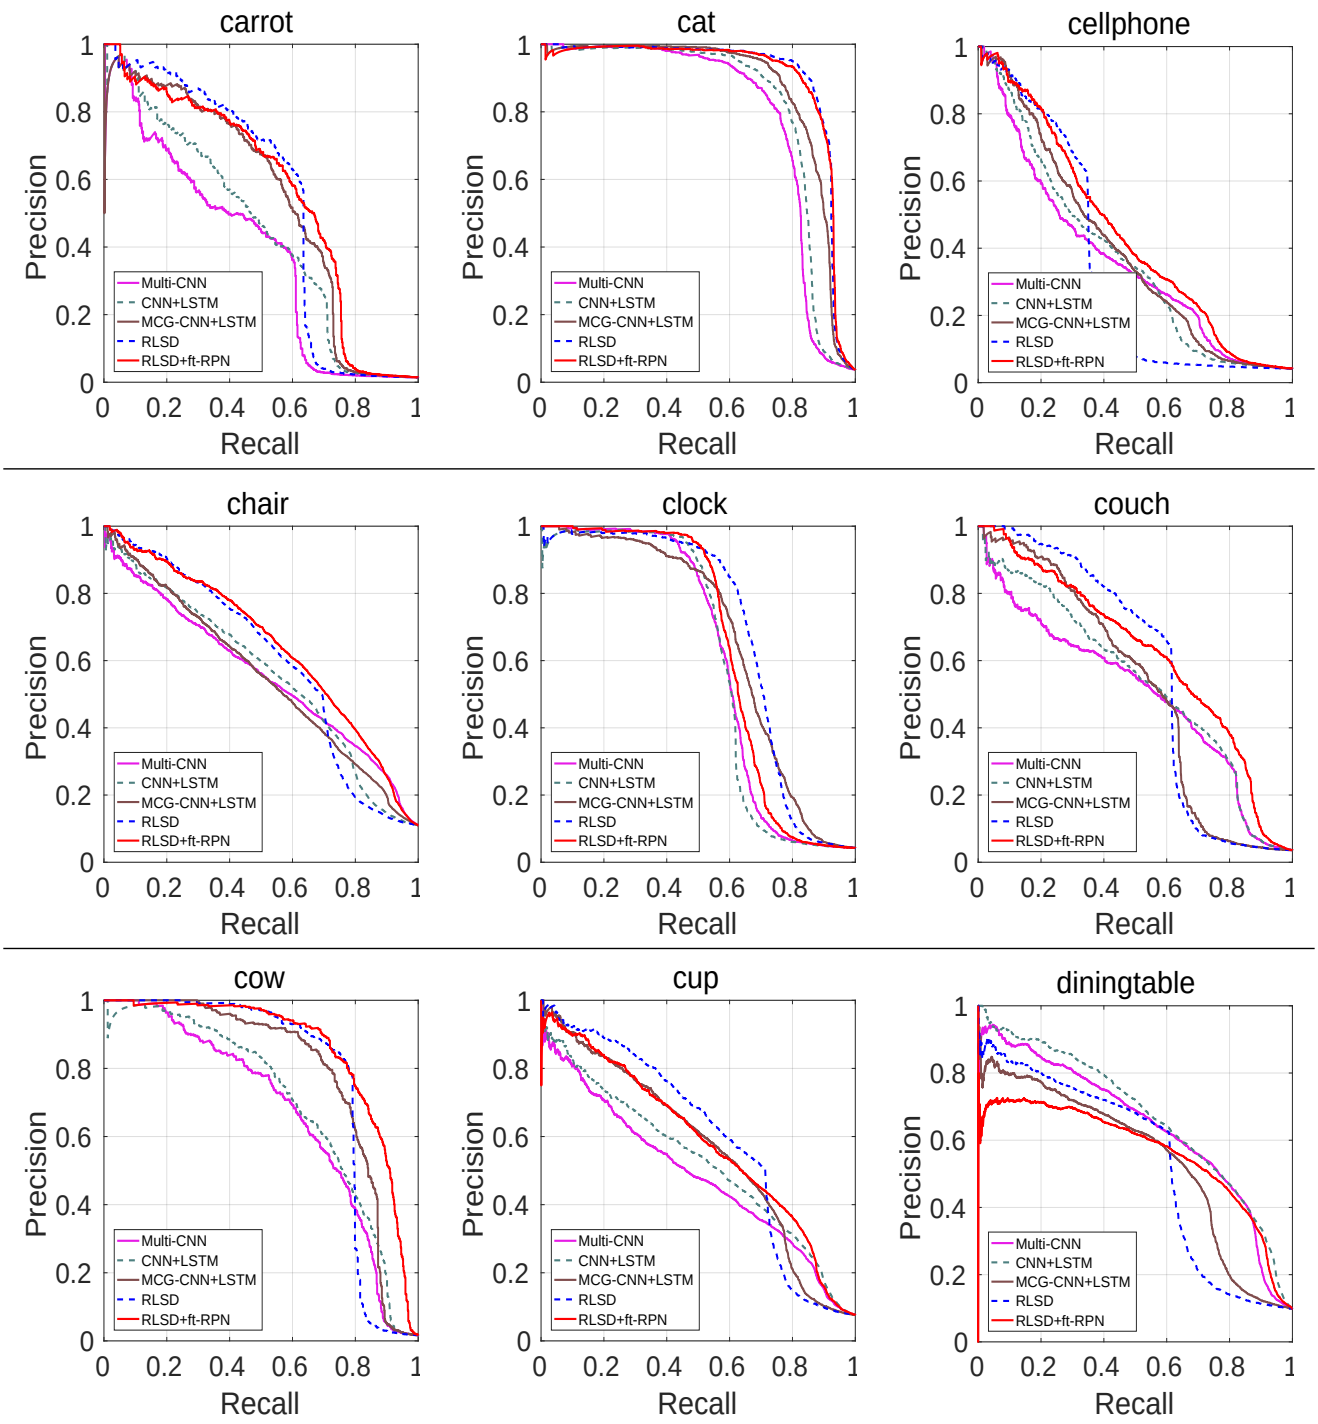

Figure 10: Precision-Recall curves for all 19th-27th classes in the MS COCO dataset, for our RLSD models and the baseline models.

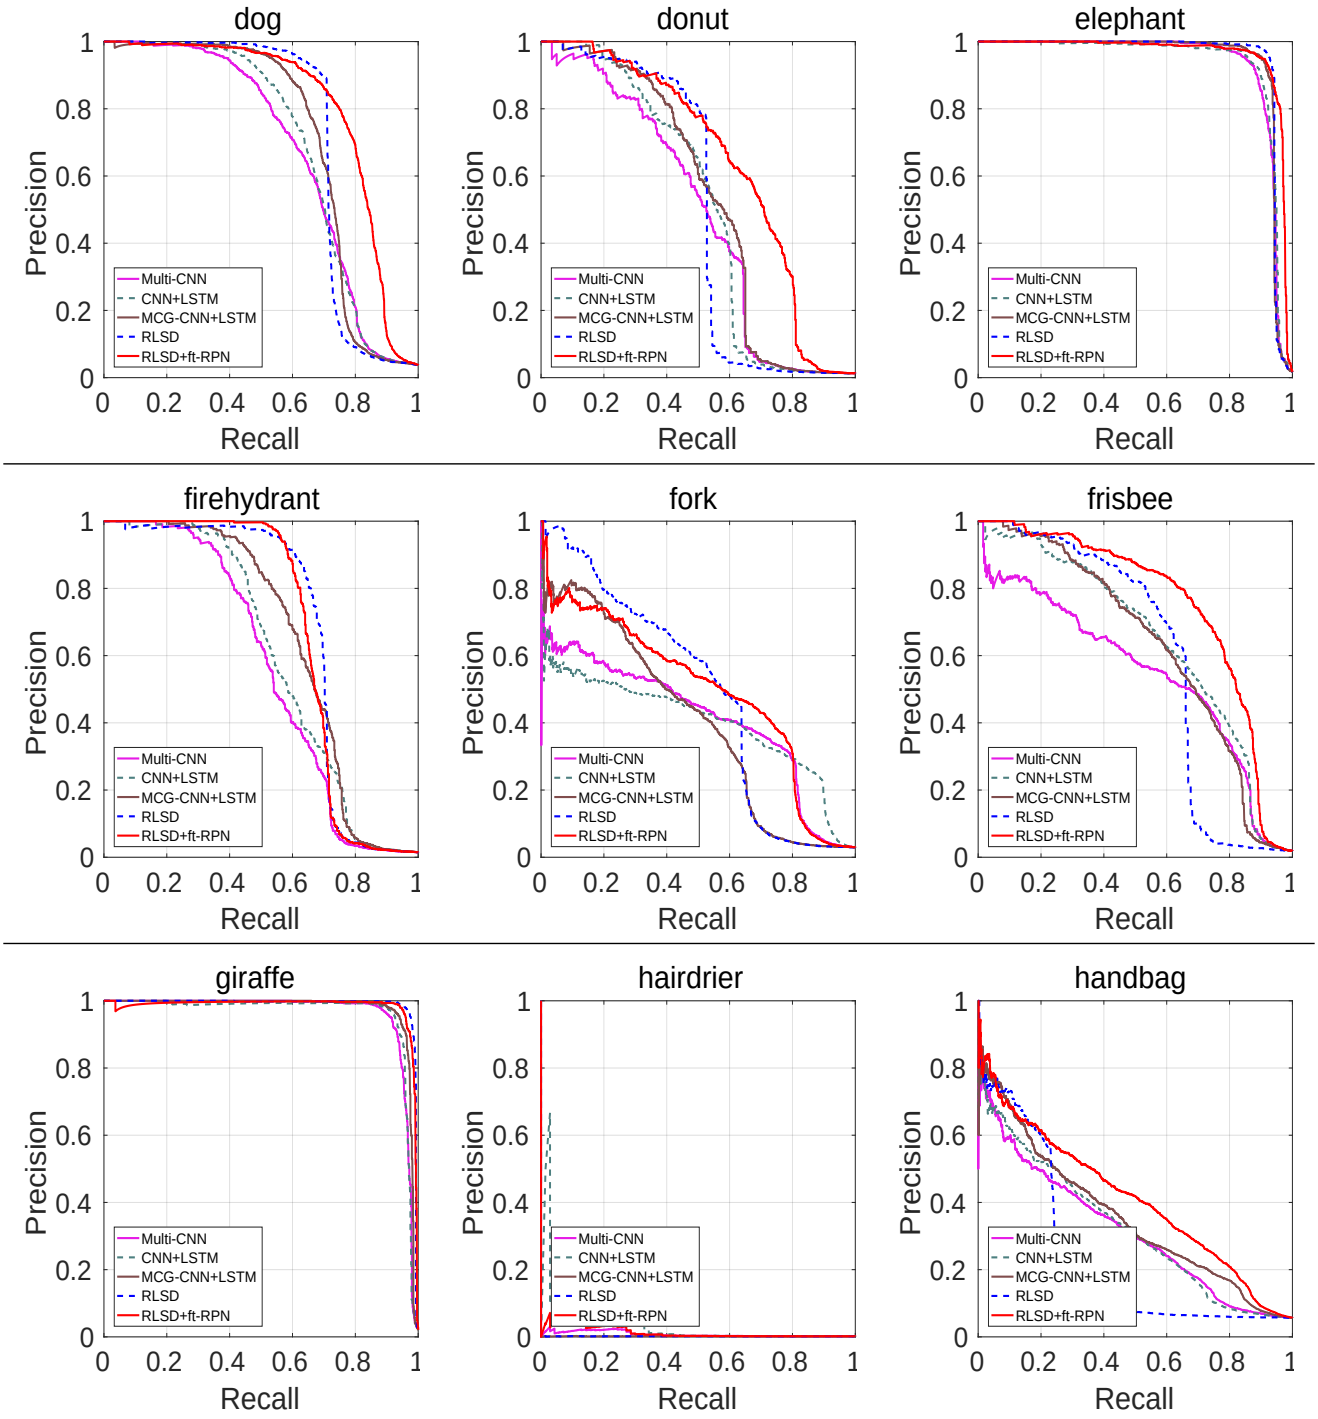

Figure 11: Precision-Recall curves for all 28th-36th classes in the MS COCO dataset, for our RLSD models and the baseline models.

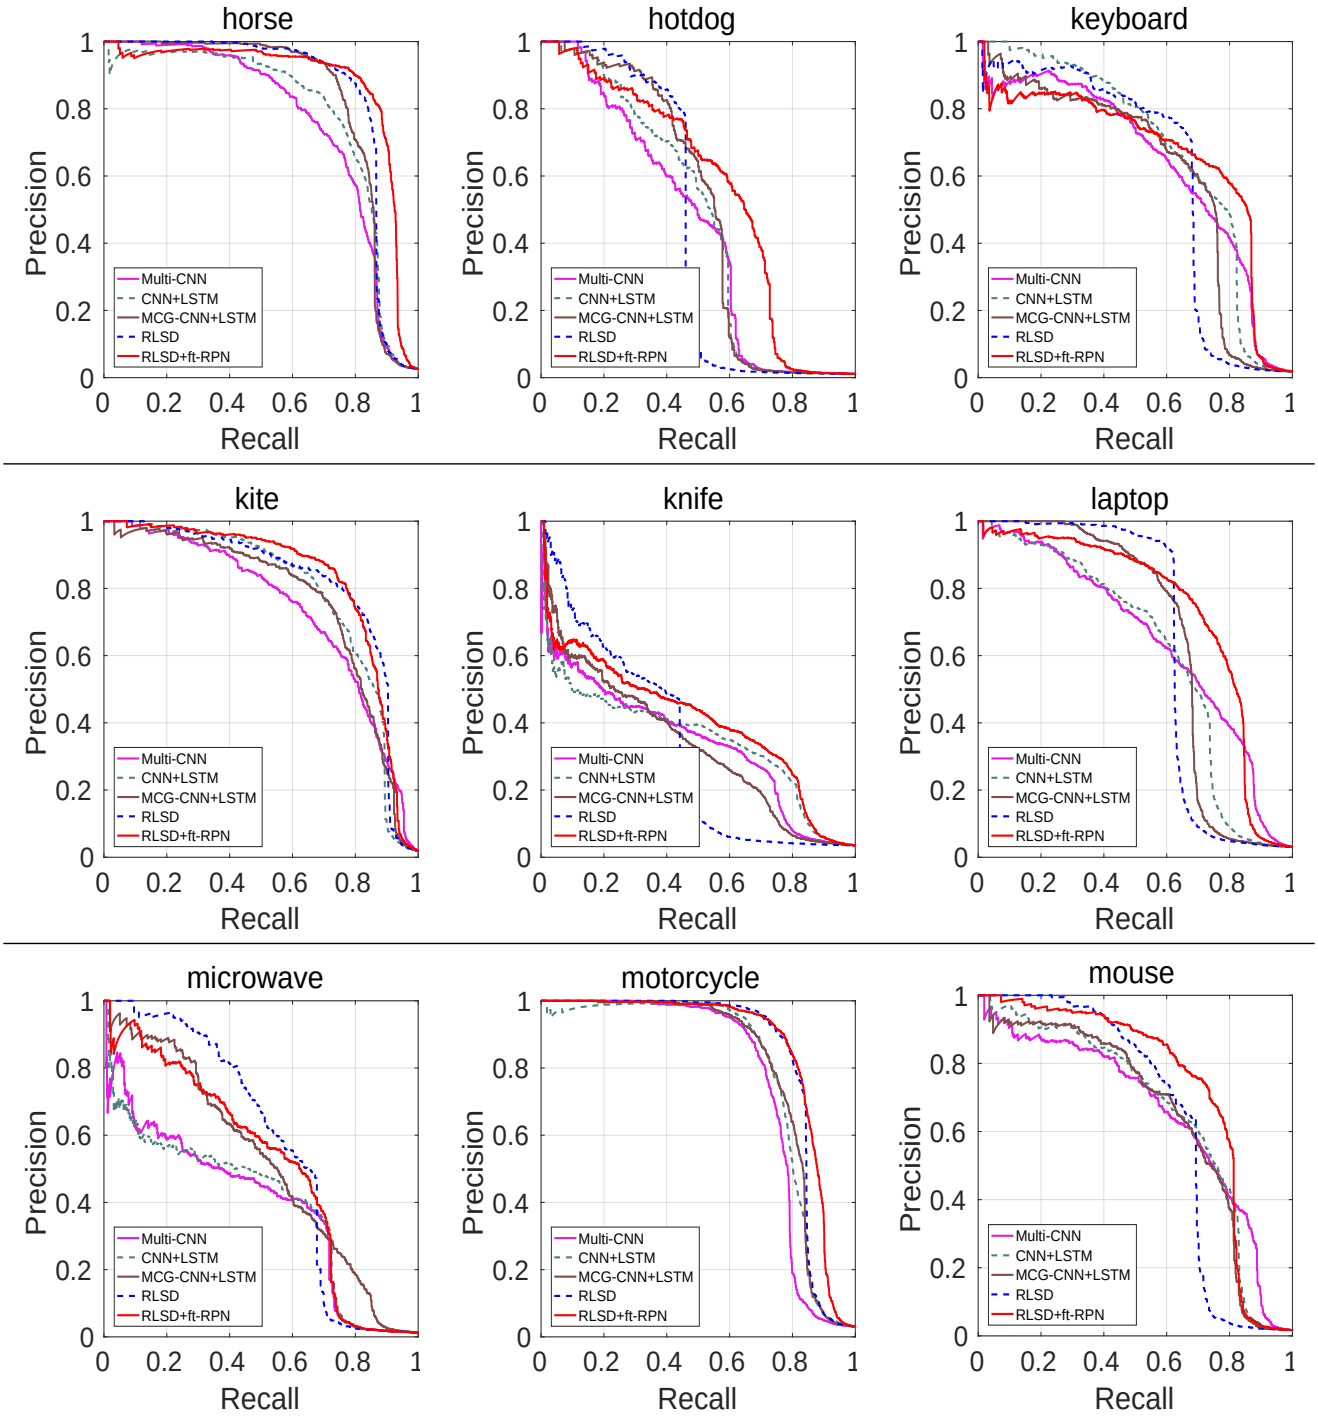

Figure 12: Precision-Recall curves for all 37th-45th classes in the MS COCO dataset, for our RLSD models and the baseline models.

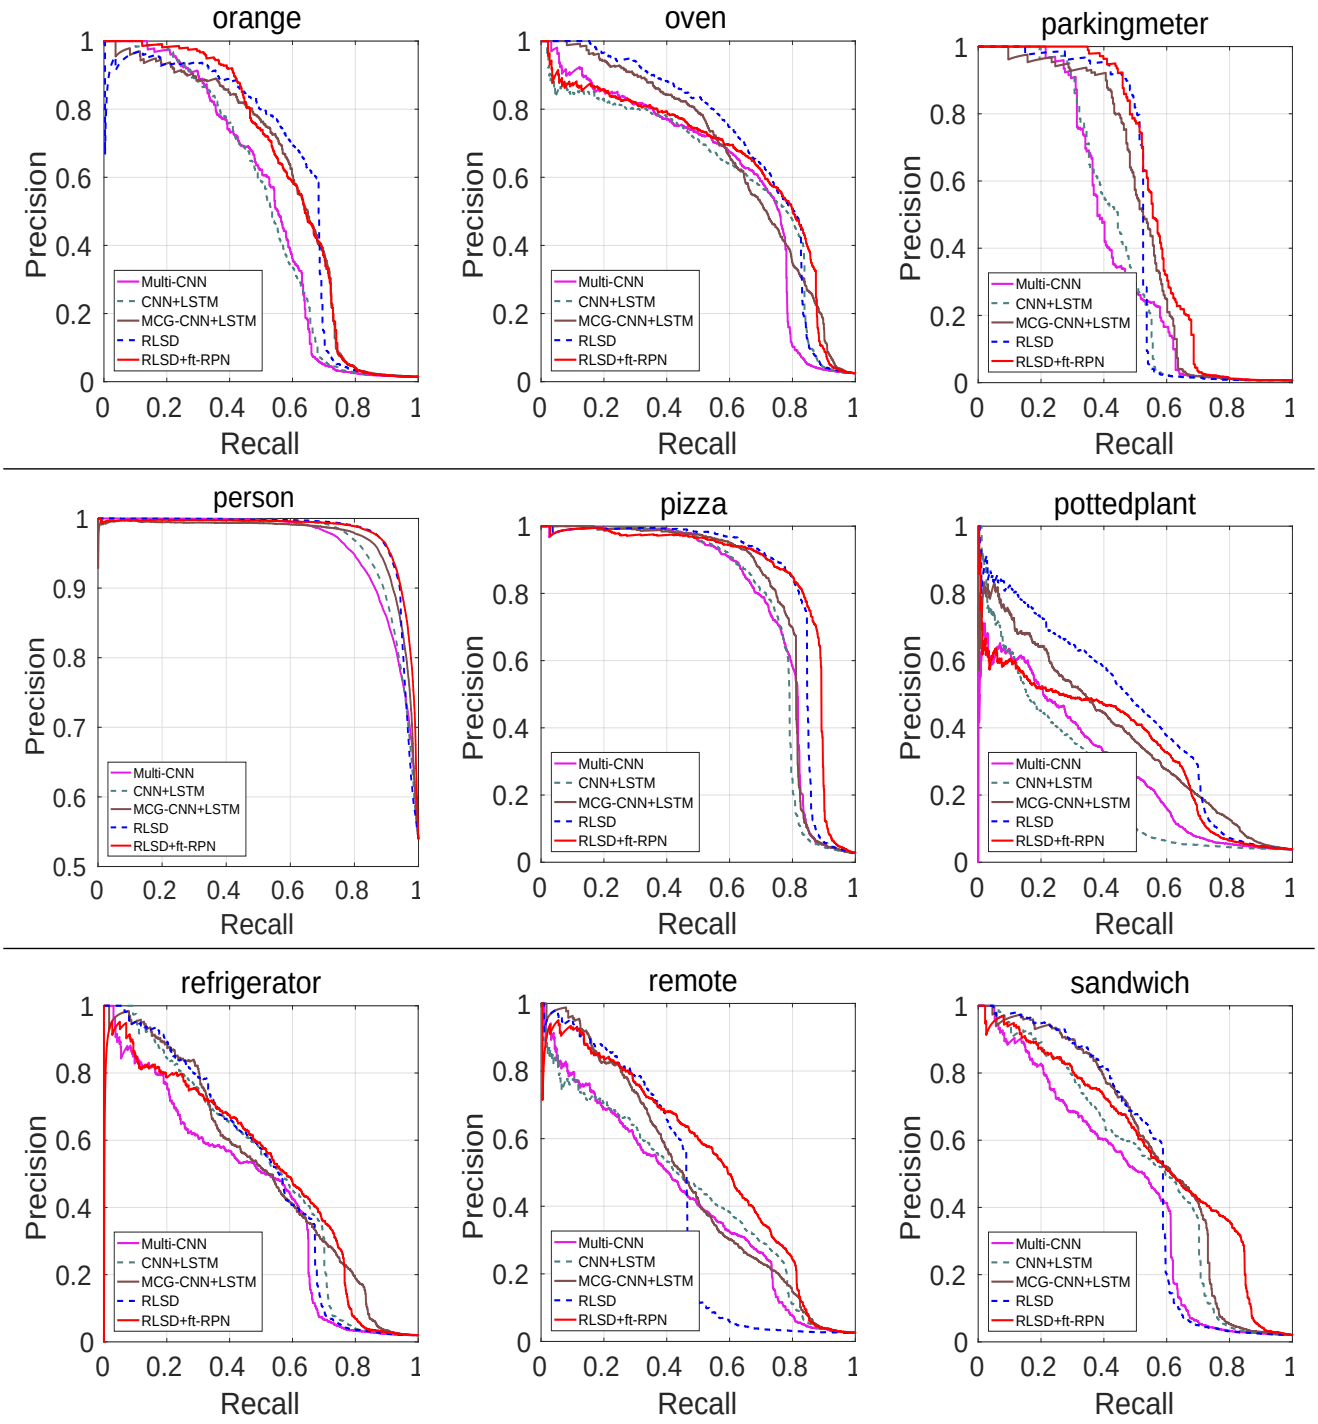

Figure 13: Precision-Recall curves for all 46th-54th classes in the MS COCO dataset, for our RLSD models and the baseline models.

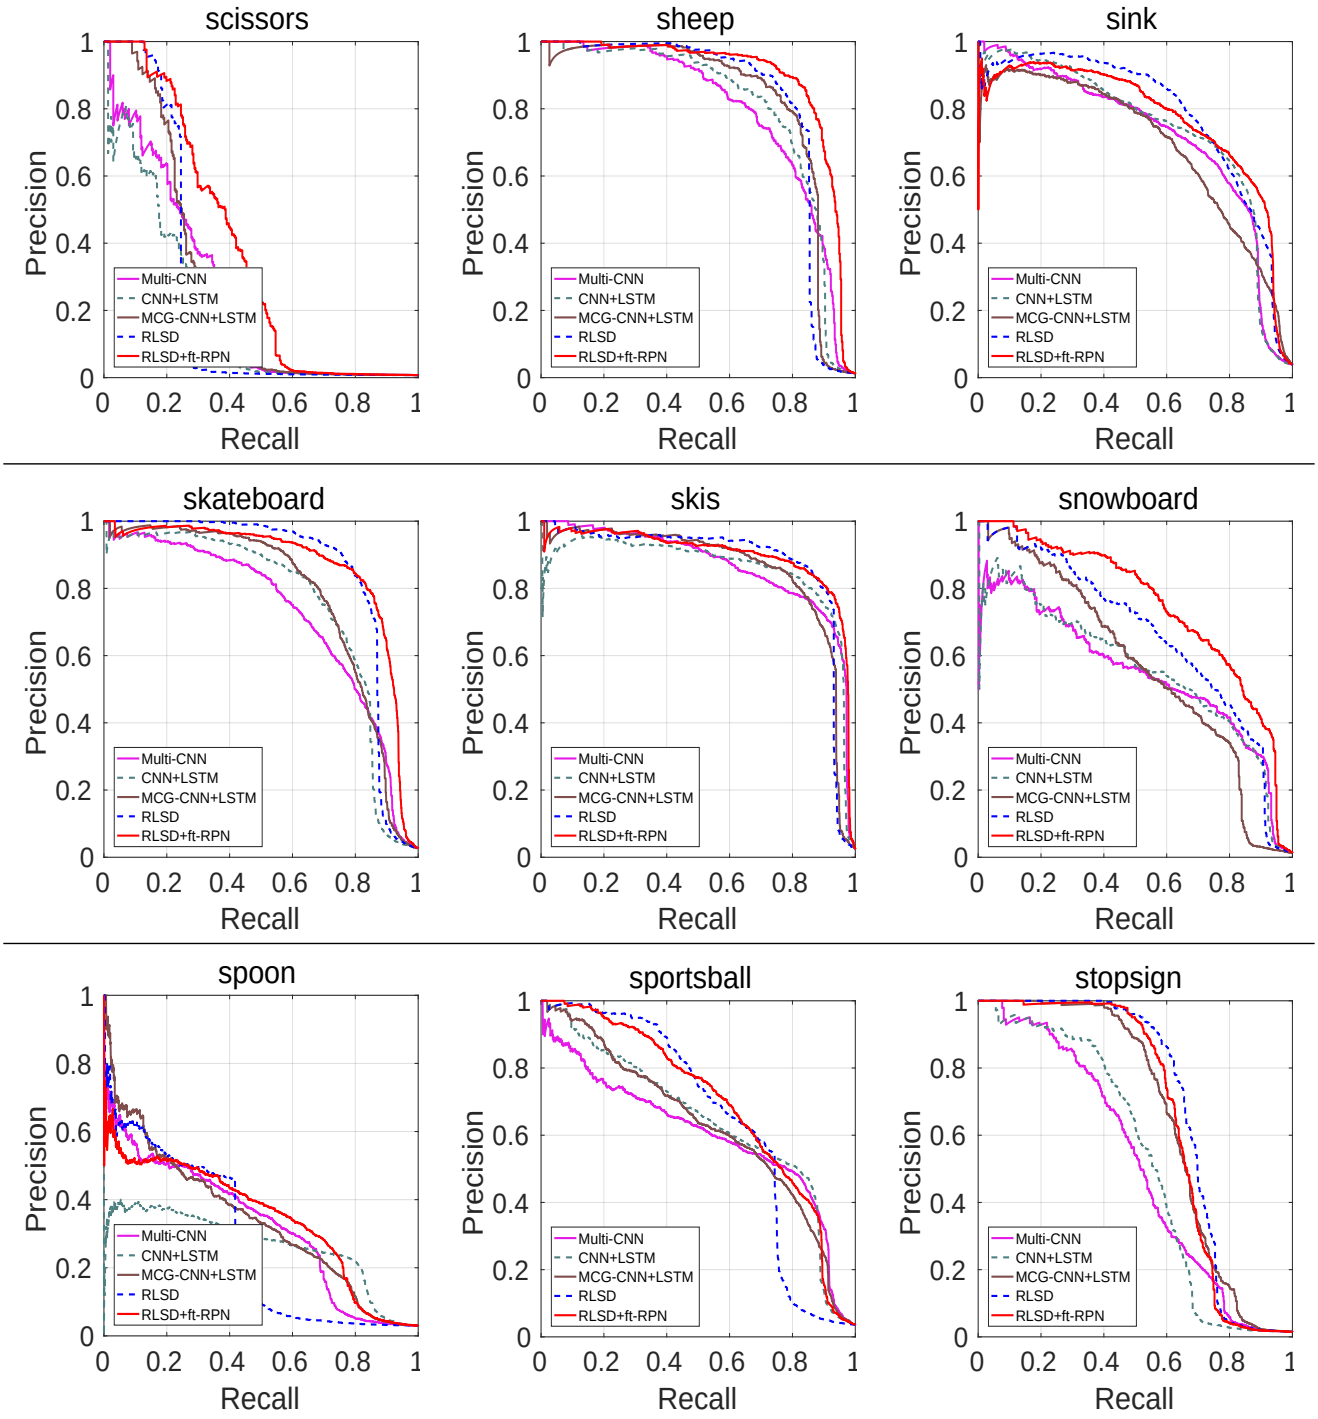

Figure 14: Precision-Recall curves for all 55th-63rd classes in the MS COCO dataset, for our RLSD models and the baseline models.

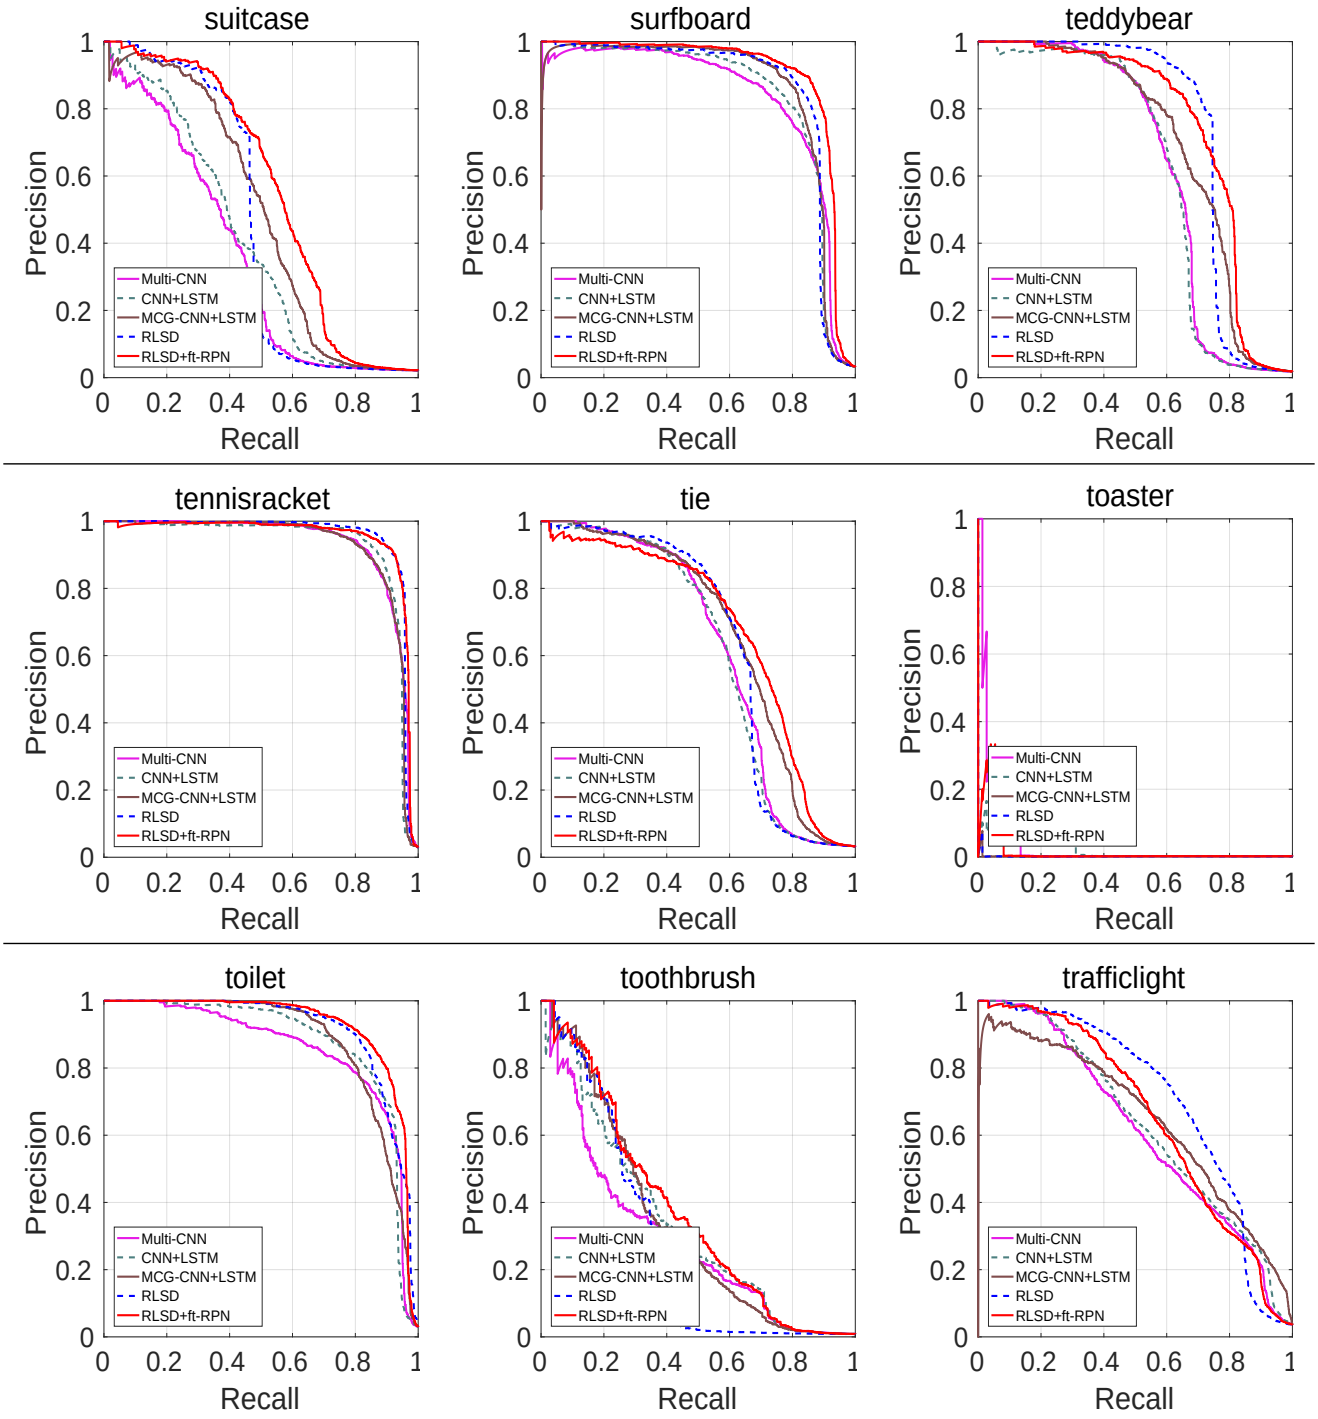

Figure 15: Precision-Recall curves for all 64th-72nd classes in the MS COCO dataset, for our RLSD models and the baseline models.

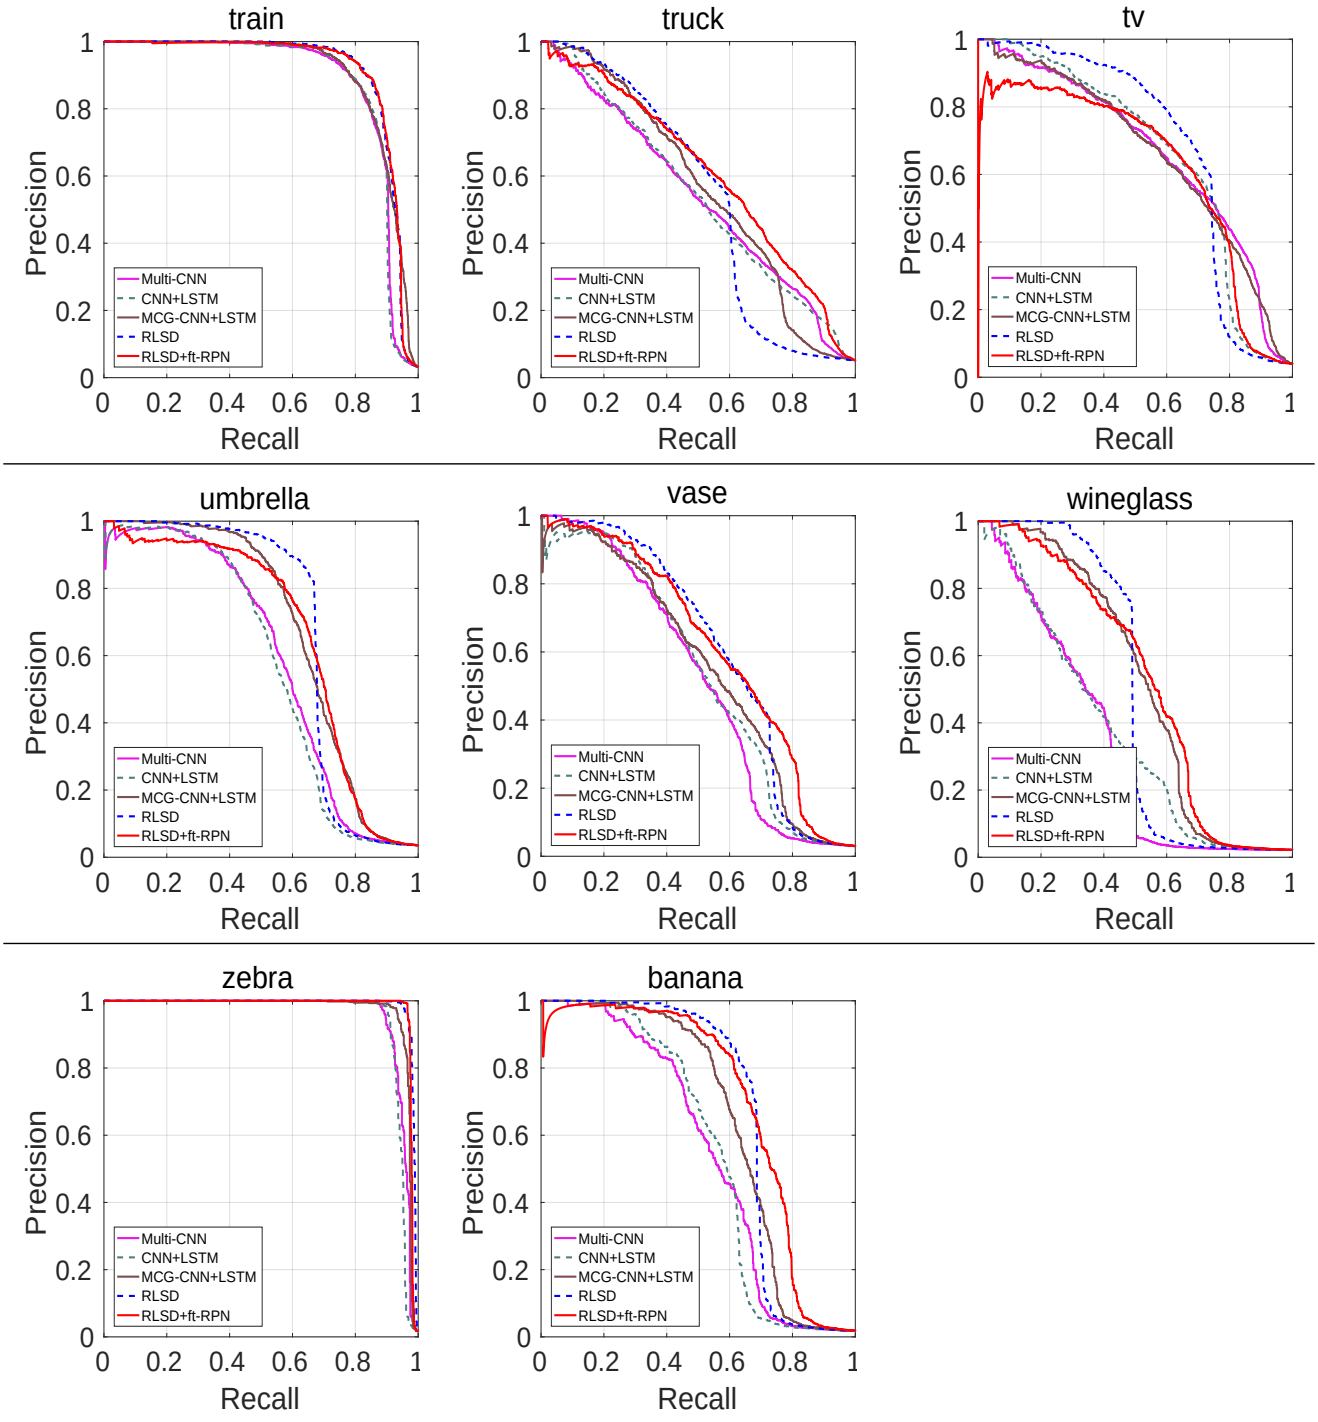

Figure 16: Precision-Recall curves for all 73rd-80th classes in the MS COCO dataset, for our RLSD models and the baseline models.

Figure 17: Some example multi-label classification results from the MS COCO, see more in the supplementary material